

Global Optimization Algorithms

– Theory and Application –



Evolutionary Algorithms.....	67
Genetic Algorithms.....	113
Genetic Programming.....	127
Learning Classifier Systems.....	197
Hill Climbing.....	209
Simulated Annealing.....	217
Example Applications.....	271
Sigoa – Implementation in Java.....	391
Background (Mathematics, Computer Science, ...)	407

Thomas Weise
Version: 2008-08-12

Newest Version: <http://www.it-weise.de/>

Preface

This e-book is devoted to global optimization algorithms, which are methods to find optimal solutions for given problems. It especially focuses on Evolutionary Computation by discussing evolutionary algorithms, genetic algorithms, Genetic Programming, Learning Classifier Systems, Evolution Strategy, Differential Evolution, Particle Swarm Optimization, and Ant Colony Optimization. It also elaborates on other meta-heuristics like Simulated Annealing, hill climbing, TS, and Random Optimization. The book is no book in the conventional sense. Because of frequent updates and changes, it is not really intended for sequential reading but more as some sort of material collection, encyclopedia, or reference work where you can look up stuff, find the correct context, and are provided with fundamentals.

With this book, two major audience groups are addressed:

1. It can help students since we try to describe the algorithms in an understandable, consistent way and, maybe even more important, includes much of the background knowledge needed to understand them. Thus, you can find summaries on stochastic theory and theoretical computer science in Part IV on page 407. Additionally, application examples are provided which give an idea how problems can be tackled with the different techniques and what results can be expected.
2. Fellow researchers and PhD students maybe may find the application examples helpful too. For them, in-depth discussions on the single methodologies are included that are supported with a large set of useful literature references.

If this book contains something you want to cite or reference in your work, please use the *citation suggestion* provided in Chapter E on page 543.

In order to maximize the utility of this electronic book, it contains automatic, clickable links. They are shaded with dark gray so the book is still b/w printable. You can click on

- entries in the table of contents,
- citation references like [1],
- page references like “67”,
- references such as “see Figure 2.1 on page 68” to sections, figures, tables, and listings, and
- URLs and links like “<http://www.lania.mx/~ccoello/EM00/> [accessed 2007-10-25]”.¹

Furthermore, we have included a list of the abbreviations and symbols used in this book in Chapter A on page 515.

¹ URLs are usually annotated with the date we have accessed them, like <http://www.lania.mx/~ccoello/EM00/> [accessed 2007-10-25]. We can neither guarantee that their content remains unchanged, nor that these sites stay available. We also assume no responsibility for anything we linked to.

The following scenario is now for example possible: A student reads the text and finds a passage that she wants to investigate in-depth. She clicks on a citation in that seems interesting and the corresponding reference is shown. To some of the references which are online available, links are provided in the reference text. By clicking on such a link, the Adobe Reader®² will open another window and load the regarding document (or a browser window of a site that links to the document). After reading it, the student may use the “backwards” button in the navigation utility to go back to the text initially read in the e-book.

The contents of this book are divided into four parts. In the first part, different optimization technologies will be introduced and their features are described. Often, small examples will be given in order to ease understanding. In the second part starting at page 271, we elaborate on different application examples in detail. With the Sigoa framework, one possible implementation of optimization algorithms in Java, is discussed and we show how some of solutions of the previous problem instances can be realized in Part III on page 391. Finally, in the last part following at page 407, the background knowledge is provided for the rest of the book. Optimization is closely related to stochastic, and hence, an introduction into this subject can be found here. Other important background information concerns theoretical computer science and clustering algorithms.

However, this book is currently worked on. It is still in a very preliminary phase where major parts are still missing or under construction. Other sections or texts are incomplete (tagged with

TODO). There may as well be errors in the contents or issues may be stated ambiguously (I do not have proof-readers). Additionally, the sequence of the content is not very good. Because of frequent updates, small sections may grow and become chapters, be moved to another place, merged with other sections, and so on. Thus, this book will change often. I choose to update, correct, and improve this book continuously instead of providing a new version each half year or so because I think this way it has a higher utility because it provides more information earlier. By doing so, I also risk confusing you with strange grammar and structure, so if you find something fishy, please let me know so I can correct and improve it right away.

The updates and improvements will result in new versions of the book, which will regularly appear on the website <http://www.it-weise.de/>. The direct download link to the newest version of this book is <http://www.it-weise.de/projects/book.pdf>. The L^AT_EX source code of this book including all graphics and the bibliography is available at <http://www.it-weise.de/projects/bookSource.zip>. The source may not always be the one of the most current version of the book. Compiling it requires multiple runs of BibT_EX because of the nifty way the references are incorporated.

I would be very happy if you provide feedback, report errors or missing things that you have found, criticize something, or have any additional ideas or suggestions. Do not hesitate to contact me via my email address tweise@gmx.de.

Copyright © 2006-2008 Thomas Weise.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled *GNU Free Documentation License (FDL)*. You can find a copy of the GNU Free Documentation Li-

² The Adobe Reader® is available for download at <http://www.adobe.com/products/reader/>
[accessed 2007-08-13].

cense in appendix Chapter B on page 527.

At many places in this book we refer to Wikipedia [2] which is a great source of knowledge. Wikipedia [2] contains articles and definitions for many of the aspects discussed in this book. Like this book, it is updated and improved frequently. Therefore, including the links adds greatly to the book's utility, in my opinion.

Important Notice

Be aware that this version of this book marks a point of transition from the first edition to the second one. Major fractions of the text of the first edition have not yet been revised and are, thus, not included in this document. However, I believe that this version corrects many shortcomings as well as inconsistencies from the first edition plus is better structured.

Contents

Preface	V
Contents.....	IX

Part I Global Optimization

1 Introduction	3
1.1 A Classification of Optimization Algorithms	4
1.1.1 Classification According to Method of Operation	4
1.1.2 Classification According to Properties	6
1.2 What is an optimum?	7
1.2.1 Single Objective Functions	7
1.2.2 Multiple Objective Functions	9
1.2.3 Weighted Sum	11
1.2.4 Pareto Optimization	13
1.2.5 The Method of Inequalities	16
1.2.6 External Decision Maker	18
1.2.7 Prevalence Optimization	19
1.3 The Structure of Optimization	21
1.3.1 Spaces, Sets, and Elements	21
1.3.2 Fitness Landscapes and Global Optimization	29
1.3.3 Gradient Descend	34
1.3.4 Other General Features	34
1.4 Problems in Optimization	37
1.4.1 Premature Convergence and Multimodality	37
1.4.2 Ruggedness and Weak Causality	40
1.4.3 Deceptiveness	43
1.4.4 Neutrality and Redundancy.....	43
1.4.5 Epistasis	47
1.4.6 Overfitting and Oversimplification	48
1.4.7 Robustness and Noise	51
1.4.8 Dynamically Changing Fitness Landscape	53
1.4.9 No Free Lunch Theorem	54
1.5 Formae and Search Space/Operator Design	55
1.5.1 Forma Analysis	56
1.5.2 Genome Design	58
1.6 General Information	60
1.6.1 Areas Of Application	60

X CONTENTS

1.6.2	Conferences, Workshops, etc.	61
1.6.3	Journals.....	64
1.6.4	Online Resources	65
1.6.5	Books.....	65
2	Evolutionary Algorithms.....	67
2.1	Introduction	67
2.1.1	The Basic Principles from Nature	67
2.1.2	Classification of Evolutionary Algorithms.....	72
2.1.3	Configuration Parameters of evolutionary algorithms.....	76
2.2	General Information	76
2.2.1	Areas Of Application	76
2.2.2	Conferences, Workshops, etc.	77
2.2.3	Journals.....	80
2.2.4	Online Resources	80
2.2.5	Books.....	81
2.3	Fitness Assignment	82
2.3.1	Introduction	82
2.3.2	Weighted Sum Fitness Assignment	83
2.3.3	Pareto Ranking	83
2.3.4	Sharing Functions	85
2.3.5	Variety Preserving Ranking.....	87
2.3.6	Tournament Fitness Assignment.....	92
2.4	Selection	92
2.4.1	Introduction	92
2.4.2	Truncation Selection	94
2.4.3	Fitness Proportionate Selection	95
2.4.4	Tournament Selection.....	100
2.4.5	Ordered Selection	103
2.4.6	Ranking Selection	105
2.4.7	VEGA Selection.....	105
2.4.8	Simple Convergence Prevention	106
2.5	Reproduction	108
2.5.1	NCGA Reproduction	109
2.6	Algorithms	110
2.6.1	VEGA	110
3	Genetic Algorithms.....	113
3.1	Introduction	113
3.2	General Information	114
3.2.1	Areas Of Application	114
3.2.2	Conferences, Workshops, etc.	115
3.2.3	Online Resources	115
3.2.4	Books.....	117
3.3	Genomes in Genetic Algorithms	117
3.4	Fixed-Length String Chromosomes	119
3.4.1	Creation	119
3.4.2	Mutation.....	119
3.4.3	Permutation	120
3.4.4	Crossover	120
3.5	Variable-Length String Chromosomes	121
3.5.1	Creation	121
3.5.2	Mutation.....	121
3.5.3	Crossover	121

3.6	Genotype-Phenotype Mappings and Artificial Embryogeny	122
3.7	Schema Theorem	122
3.7.1	Schemata and Masks	123
3.7.2	Wildcards	123
3.7.3	Holland's Schema Theorem	124
3.7.4	Criticism of the Schema Theorem	124
3.7.5	The Building Block Hypothesis	125
4	Genetic Programming	127
4.1	Introduction	127
4.1.1	History	127
4.2	General Information	129
4.2.1	Areas Of Application	129
4.2.2	Conferences, Workshops, etc.	130
4.2.3	Journals	131
4.2.4	Online Resources	131
4.2.5	Books	132
4.3	(Standard) Tree Genomes	132
4.3.1	Creation	132
4.3.2	Mutation	133
4.3.3	Recombination	133
4.3.4	Permutation	135
4.3.5	Editing	135
4.3.6	Encapsulation	136
4.3.7	Wrapping	136
4.3.8	Lifting	137
4.3.9	Automatically Defined Functions	137
4.3.10	Automatically Defined Macros	138
4.3.11	Node Selection	139
4.4	Genotype-Phenotype Mappings	141
4.4.1	Cramer's Genetic Programming	141
4.4.2	Binary Genetic Programming	142
4.4.3	Gene Expression Programming	142
4.4.4	Edge Encoding	144
4.5	Grammars in Genetic Programming	147
4.5.1	Introduction	147
4.5.2	Trivial Approach	147
4.5.3	Strongly Typed Genetic Programming	148
4.5.4	Early Research in GGGP	149
4.5.5	Gads 1	149
4.5.6	Grammatical Evolution	151
4.5.7	Gads 2	155
4.5.8	Christiansen Grammar Evolution	156
4.5.9	Tree-Adjoining Grammar-guided Genetic Programming	158
4.6	Linear Genetic Programming	161
4.6.1	Homologous Crossover	163
4.6.2	Page-based LGP	163
4.7	Graph-based Approaches	163
4.7.1	Parallel Algorithm Discovery and Orchestration	163
4.7.2	Parallel Distributed Genetic Programming	164
4.7.3	Genetic Network Programming	166
4.7.4	Cartesian Genetic Programming	167
4.8	Epistasis in Genetic Programming	169
4.8.1	Introduction	169

XII CONTENTS

4.8.2	Algorithmic Chemistry	171
4.8.3	Rule-based Genetic Programming	173
4.8.4	Soft Assignment	179
4.9	Artificial Life and Artificial Chemistry	180
4.9.1	Push, PushGP, and Pushpop	180
4.9.2	Fraglets	183
4.10	Evolving Algorithms	186
4.10.1	Restricting Problems	186
4.10.2	Why No Exhaustive Testing?	187
4.10.3	All-Or-Nothing?	188
4.10.4	Non-Functional Features of Algorithms	188
5	Evolution Strategy	191
5.1	Introduction	191
5.2	General Information	191
5.2.1	Areas Of Application	191
5.2.2	Conferences, Workshops, etc.	192
5.2.3	Books	192
5.3	Populations in Evolution Strategy	192
5.3.1	$(1 + 1)$ -ES	192
5.3.2	$(\mu + 1)$ -ES	192
5.3.3	$(\mu + \lambda)$ -ES	192
5.3.4	(μ, λ) -ES	193
5.3.5	$(\mu/\rho, \lambda)$ -ES	193
5.3.6	$(\mu/\rho + \lambda)$ -ES	193
5.3.7	$(\mu', \lambda'(\mu, \lambda)^\gamma)$ -ES	193
5.4	One-Fifth Rule	193
5.5	Differential Evolution	193
5.5.1	Introduction	193
5.5.2	General Information	194
6	Evolutionary Programming	195
6.1	Introduction	195
6.2	General Information	195
6.2.1	Areas Of Application	195
6.2.2	Conferences, Workshops, etc.	196
6.2.3	Books	196
7	Learning Classifier Systems	197
7.1	Introduction	197
7.2	General Information	197
7.2.1	Areas Of Application	197
7.2.2	Conferences, Workshops, etc.	197
7.2.3	Books	198
7.3	The Basic Idea of Learning Classifier Systems	198
7.3.1	Messages	198
7.3.2	Conditions	200
7.3.3	Actions	202
7.3.4	Classifiers	202
7.3.5	Non-Learning Classifier Systems	203
7.3.6	Learning Classifier Systems	204
7.3.7	The Bucket Brigade Algorithm	204
7.3.8	Applying the Genetic Algorithm	206
7.4	Families of Learning Classifier Systems	206

8	Hill Climbing	209
8.1	Introduction	209
8.2	General Information	210
8.2.1	Areas Of Application	210
8.3	Multi-Objective Hill Climbing	210
8.4	Problems in Hill Climbing	211
8.5	Hill Climbing with Random Restarts	212
9	Random Optimization	213
9.1	Introduction	213
9.2	General Information	214
9.2.1	Areas Of Application	214
10	Simulated Annealing	217
10.1	Introduction	217
10.2	General Information	219
10.2.1	Areas Of Application	219
10.2.2	Books	219
10.3	Temperature Scheduling	219
10.4	Multi-Objective Simulated Annealing	220
11	Downhill Simplex (Nelder and Mead)	223
11.1	Introduction	223
11.2	General Information	223
11.2.1	Areas Of Application	223
11.2.2	Online Resources	224
11.2.3	Books	224
11.3	The Downhill Simplex Algorithm	224
11.4	Hybridizing with the Downhill Simplex	226
12	Tabu Search	229
12.1	Introduction	229
12.2	General Information	230
12.2.1	Areas Of Application	230
12.2.2	Books	230
12.3	Multi-Objective Tabu Search	230
13	Ant Colony Optimization	233
13.1	Introduction	233
13.2	General Information	234
13.2.1	Areas Of Application	234
13.2.2	Conferences, Workshops, etc.	234
13.2.3	Journals	234
13.2.4	Online Resources	234
14	Particle Swarm Optimization	237
14.1	Introduction	237
14.2	General Information	238
14.2.1	Areas Of Application	238
14.2.2	Online Resources	239
14.2.3	Conferences, Workshops, etc.	239
14.2.4	Books	239

XIV CONTENTS

15	Memetic Algorithms	241
15.1	General Information	242
15.1.1	Areas Of Application	242
15.1.2	Online Resources	242
15.1.3	Books	243
16	State Space Search	245
16.1	Introduction	245
16.2	General Information	247
16.2.1	Areas Of Application	247
16.2.2	Books	247
16.3	Uninformed Search	247
16.3.1	Breadth-First Search	247
16.3.2	Depth-First Search	248
16.3.3	Depth-limited Search	249
16.3.4	Iterative Deepening Depth-First Search	249
16.3.5	Random Walks	250
16.4	Informed Search	251
16.4.1	Greedy Search	251
16.4.2	A* search	252
16.4.3	Adaptive Walks	252
17	Parallelization and Distribution	255
17.1	Analysis	255
17.2	Distribution	257
17.2.1	Client-Server	257
17.2.2	Island Model	258
17.2.3	Mixed Distribution	261
17.3	Cellular Genetic Algorithms	261
18	Maintaining the Optimal Set	263
18.1	Updating the Optimal Set	263
18.2	Obtaining Optimal Elements	264
18.3	Pruning the Optimal Set	265
18.3.1	Pruning via Clustering	266
18.3.2	Adaptive Grid Archiving	266

Part II Applications

19	Experimental Settings, Measures, and Evaluations	271
19.1	Settings	271
19.1.1	The Optimization Problem	271
19.1.2	The Optimization Algorithm Applied	272
19.1.3	Other Run Parameters	273
19.2	Measures	274
19.3	Evaluation	275
19.3.1	Simple Evaluation Measures	275
19.3.2	Sophisticated Estimates	277

20	Benchmarks and Toy Problems	279
20.1	Real Problem Spaces	279
20.1.1	Single-Objective Optimization	279
20.1.2	Multi-Objective Optimization	280
20.1.3	Dynamic Fitness Landscapes	280
20.2	Binary Problem Spaces	281
20.2.1	Kauffman's NK Fitness Landscapes	281
20.2.2	The p-Spin Model	284
20.2.3	The ND Family of Fitness Landscapes	285
20.2.4	The Royal Road	285
20.2.5	OneMax and BinInt	289
20.2.6	Tunable Model for Problematic Phenomena	290
20.3	Genetic Programming Problems	303
20.3.1	Artificial Ant	303
20.3.2	The Greatest Common Divisor	305
21	Contests	317
21.1	DATA-MINING-CUP	317
21.1.1	Introduction	317
21.1.2	The 2007 Contest – Using Classifier Systems	318
21.2	The Web Service Challenge	327
21.2.1	Introduction	327
21.2.2	The 2006/2007 Semantic Challenge	329
22	Real-World Applications	341
22.1	Symbolic Regression	341
22.1.1	Genetic Programming: Genome for Symbolic Regression	341
22.1.2	Sample Data, Quality, and Estimation Theory	342
22.1.3	An Example and the Phenomenon of Overfitting	343
22.1.4	Limits of Symbolic Regression	345
23	Research Applications	347
23.1	Global Optimization of Distributed Systems	347
23.1.1	Optimizing Network Topology and Dimensioning	348
23.1.2	Optimizing Routing	352
23.1.3	Synthesizing Protocols	359
23.1.4	Optimizing Network Security	363
23.1.5	Optimizing Parameters and Configurations	364
23.2	Genetic Programming of Distributed Algorithms	365
23.2.1	Introduction	365
23.2.2	Evolving Proactive Aggregation Protocols	366

Part III Sigoa – Implementation in Java

24	Introduction	391
24.1	Requirements Analysis	392
24.1.1	Multi-Objectivity	392
24.1.2	Separation of Specification and Implementation	392
24.1.3	Separation of Concerns	393
24.1.4	Support for Pluggable Simulations and Introspection	393
24.1.5	Distribution utilities	393
24.2	Architecture	393
24.3	Subsystems	395

XVI CONTENTS

25 Examples	397
25.1 The 2007 DATA-MINING-CUP	397
25.1.1 The Phenotype	397
25.1.2 The Genotype and the Embryogeny	398
25.1.3 The Simulation	398
25.1.4 The Objective Functions	401
25.1.5 The Evolution Process	403

Part IV Background

26 Set Theory	407
26.1 Set Membership	407
26.2 Relations between Sets	407
26.3 Special Sets	408
26.4 Operations on Sets	408
26.5 Tuples	410
26.6 Lists	411
26.7 Binary Relations	413
26.7.1 Functions	414
26.7.2 Order Relations	415
26.7.3 Equivalence Relations	416
27 Stochastic Theory and Statistics	417
27.1 Probability	417
27.1.1 Probably as defined by Bernoulli (1713)	418
27.1.2 The Limiting Frequency Theory of von Mises	419
27.1.3 The Axioms of Kolmogorov	419
27.1.4 Conditional Probability	420
27.1.5 Random Variable	421
27.1.6 Cumulative Distribution Function	421
27.1.7 Probability Mass Function	422
27.1.8 Probability Density Function	422
27.2 Stochastic Properties	423
27.2.1 Count, Min, Max and Range	423
27.2.2 Expected Value and Arithmetic Mean	424
27.2.3 Variance and Standard Deviation	424
27.2.4 Moments	426
27.2.5 Skewness and Kurtosis	427
27.2.6 Median, Quantiles, and Mode	427
27.2.7 Entropy	429
27.2.8 The Law of Large Numbers	429
27.3 Some Discrete Distributions	429
27.3.1 Discrete Uniform Distribution	430
27.3.2 Poisson Distribution π_λ	431
27.3.3 Binomial Distribution $B(n, p)$	434
27.4 Some Continuous Distributions	435
27.4.1 Continuous Uniform Distribution	436
27.4.2 Normal Distribution $N(\mu, \sigma^2)$	437
27.4.3 Exponential Distribution $\exp(\lambda)$	440
27.4.4 Chi-square Distribution	441
27.4.5 Student's t-Distribution	445
27.5 Example – Throwing a Dice	448
27.6 Estimation Theory	450

27.6.1	Likelihood and Maximum Likelihood Estimators	451
27.6.2	Best Linear Unbiased Estimators	454
27.6.3	Confidence Intervals	454
27.6.4	Density Estimation	457
27.7	Generating Random Numbers	459
27.7.1	Generating Pseudorandom Numbers	460
27.7.2	Random Functions	461
27.7.3	Converting Random Numbers to other Distributions	462
27.8	List of Functions	466
27.8.1	Gamma Function	466
28	Clustering	467
28.1	Distance Measures	469
28.1.1	Distance Measures for Strings of Equal Length	469
28.1.2	Distance Measures for Real-Valued Vectors	469
28.1.3	Elements Representing a Cluster	471
28.1.4	Distance Measures Between Clusters	471
28.2	Clustering Algorithms	471
28.2.1	Cluster Error	471
28.2.2	k -means Clustering	472
28.2.3	n^{th} Nearest Neighbor Clustering	472
28.2.4	Linkage Clustering	473
28.2.5	Leader Clustering	475
29	Theoretical Computer Science	479
29.1	Introduction	479
29.1.1	Algorithms and Programs	479
29.1.2	Properties of Algorithms	481
29.1.3	Complexity of Algorithms	482
29.1.4	Randomized Algorithms	484
29.2	Distributed Systems and Distributed Algorithms	485
29.2.1	Network Topologies	486
29.2.2	Some Architectures of Distributed Systems	488
29.2.3	Modeling Distributed Systems	493
29.3	Grammars and Languages	501
29.3.1	Syntax and Formal Languages	502
29.3.2	Generative Grammars	502
29.3.3	Derivation Trees	503
29.3.4	Backus-Naur Form	504
29.3.5	Extended Backus-Naur Form	505
29.3.6	Attribute Grammars	505
29.3.7	Extended Attribute Grammars	507
29.3.8	Adaptive Grammars	508
29.3.9	Christiansen Grammars	509
29.3.10	Tree-Adjoining Grammars	510
29.3.11	S-expressions	512

Part V Appendices

A	Abbreviations and Symbols	515

XVIII CONTENTS

B	GNU <u>F</u>ree <u>D</u>ocumentation <u>L</u>icense (FDL)	527
B.1	Preamble	527
B.2	Applicability and Definitions	527
B.3	Verbatim Copying	528
B.4	Copying in Quantity	529
B.5	Modifications	529
B.6	Combining Documents	531
B.7	Collections of Documents	531
B.8	Aggregation with Independent Works	531
B.9	Translation	531
B.10	Termination	532
B.11	Future Revisions of this License	532
C	GNU <u>L</u>esser <u>G</u>eneral <u>P</u>ublic <u>L</u>icense (LPG<u>L</u>)	533
C.1	Preamble	533
C.2	Terms and Conditions for Copying, Distribution and Modification	534
C.3	No Warranty	538
C.4	How to Apply These Terms to Your New Libraries	539
D	Credits and Contributors	541
E	Citation Suggestion	543
	References	545
	Index	717
	List of Figures	733
	List of Tables	739
	List of Algorithms	741
	List of Listings	743

Global Optimization

Introduction

*“I will not stop runnin’ ’till I’m winning the race
I can’t live with myself being in second place”*

— R.A. The Rugged Man [3]

One of the most fundamental principles in our world is the search for an optimal state. It begins in the microcosm where atoms in physics try to form bonds¹ in order to minimize the energy of their electrons [4]. When molecules form solid bodies during the process of freezing, they try to assume energy-optimal crystal structures. These processes, of course, are not driven by any higher intention but purely result from the laws of physics.

The same goes for the biological principle of *survival of the fittest* [5] which, together with the biological evolution [6], leads to better adaptation of the species to their environment. Here, a local optimum is a well-adapted species that dominates all other animals in its surroundings. Homo sapiens have reached this level, sharing it with ants, bacteria, flies, cockroaches, and all sorts of other creepy creatures.

As long as humankind exists, we strive for perfection in many areas. We want to reach a maximum degree of happiness with the least amount of effort. In our economy, profit and sales must be maximized and costs should be as low as possible. Therefore, optimization is one of the oldest of sciences which even extends into daily life [7].

If something is important, general, and abstract enough, there is always a mathematical discipline dealing with it. Global optimization² is the branch of applied mathematics and numerical analysis that focuses on, well, optimization. The goal of global optimization is to find the best possible elements $\mathbf{x}^*$ from a set $\mathbb{X}$ according to a set of criteria $F = \{f_1, f_2, \dots, f_n\}$. These criteria are expressed as mathematical functions³, the so-called objective functions.

Definition 1.1 (Objective Function). An objective function $f : \mathbb{X} \mapsto Y$ with $Y \subseteq \mathbb{R}$ is a mathematical function which is subject to optimization.

The codomain Y of an objective function as well as its range must be a subset of the real numbers ($Y \subseteq \mathbb{R}$). The domain $\mathbb{X}$ of f is called problem space and can represent any type of elements like numbers, lists, construction plans, and so on. It is chosen according to the problem to be solved with the optimization process. Objective functions are not necessarily mere mathematical expressions, but can be complex algorithms that, for example, involve multiple simulations. Global optimization comprises all techniques that can be used to find the best elements $\mathbf{x}^*$ in $\mathbb{X}$ with respect to such criteria $f \in F$.

In the remaining text of this introduction, we will first provide a rough classification of the different optimization techniques which we will investigate in the further course of this book (Section 1.1). In Section 1.2, we will outline how these *best* elements which we are after can be defined. We will use Section 1.3 to shed some more light onto the meaning and inter-relation of the symbols already mentioned ($f, F, x, \mathbf{x}^*, \mathbb{X}, Y, \dots$) and outline

¹ http://en.wikipedia.org/wiki/Chemical_bond [accessed 2007-07-12]

² http://en.wikipedia.org/wiki/Global_optimization [accessed 2007-07-03]

³ The concept of mathematical functions is outlined in set theory in Definition 26.27 on page 414.

the general structure of optimization processes. If optimization was a simple thing to do, there wouldn't be a whole branch of mathematics with lots of cunning people dealing with it. In Section 1.4 we will introduce the major problems that can be encountered during optimization. We will discuss Formae as a general way of describing properties of possible solutions in Section 1.5. In this book, we will provide additional hints that point to useful literature, web links, conferences, and so on for all algorithms which we discuss. The first of these information records, dealing with global optimization in general, can be found in Section 1.6.

In the chapters to follow these introductory sections, different approaches to optimization are discussed, examples for the applications are given, and the mathematical foundation and background information is provided.

1.1 A Classification of Optimization Algorithms

In this book, we will only be able to discuss a small fraction of the wide variety of global optimization techniques [8]. Before digging any deeper into the matter, I will attempt to provide a classification of these algorithms as overview and discuss some basic use cases.

1.1.1 Classification According to Method of Operation

Figure 1.1 sketches a rough taxonomy of global optimization methods. Generally, optimization algorithms can be divided in two basic classes: deterministic and probabilistic algorithms. Deterministic algorithms (see also Definition 29.11 on page 482) are most often used if a clear relation between the characteristics of the possible solutions and their utility for a given problem exists. Then, the search space can efficiently be explored using for example a divide and conquer scheme⁴. If the relation between a solution candidate and its "fitness" are not so obvious or too complicated, or the dimensionality of the search space is very high, it becomes harder to solve a problem deterministically. Trying it would possibly result in exhaustive enumeration of the search space, which is not feasible even for relatively small problems.

Then, probabilistic algorithms⁵ come into play. Especially relevant in this context are Monte Carlo⁶-based approaches. They trade in guaranteed correctness of the solution for a shorter runtime. This does not mean that the results obtained using them are incorrect – they may just not be the global optima. On the other hand, a solution a little bit inferior to the best possible one is better than one which needs 10^{100} years to be found. . .

Heuristics used in global optimization are functions that help decide which one of a set of possible solutions is to be examined next. On one hand, deterministic algorithms usually employ heuristics in order to define the processing order of the solution candidates. An example for such a strategy is informed search, as discussed in Section 16.4 on page 251. Probabilistic methods, on the other hand, may only consider those elements of the search space in further computations that have been selected by the heuristic.

Definition 1.2 (Heuristic). A heuristic⁷ [9, 10, 11] is a part of an optimization algorithm that uses the information currently gathered by the algorithm to help to decide which solution candidate should be tested next or how the next individual can be produced. Heuristics are usually problem class dependent.

⁴ http://en.wikipedia.org/wiki/Divide_and_conquer_algorithm [accessed 2007-07-09]

⁵ The common properties of probabilistic algorithms are specified in Definition 29.18 on page 484.

⁶ See Definition 29.20 on page 484 for a in-depth discussion of the Monte Carlo-type probabilistic algorithms

⁷ http://en.wikipedia.org/wiki/Heuristic_%28computer_science%29 [accessed 2007-07-03]




Fig. 1.1: The taxonomy of global optimization algorithms.

Definition 1.3 (Metaheuristic). A metaheuristic⁸ is a heuristic method for solving a very general class of problems. It combines (problem dependent) objective functions or heuristics in an abstract and hopefully efficient way.

This combination is often performed statistically by using population as sample from the search space or based on a model of some natural phenomenon or physical process. Simulated annealing for example decides which solution candidate to be evaluated according to the Boltzmann probability factor of atom configurations of solidifying metal melts. Evolutionary algorithms copy the behavior of natural evolution and treat solution candidates as individuals that compete in a virtual environment.

⁸ <http://en.wikipedia.org/wiki/Metaheuristic> [accessed 2007-07-03]

In principle, all the probabilistic optimization algorithms that we consider in this book as well as some of the deterministic ones (the greedy state space search, for instance) are metaheuristics.

An important class of probabilistic, Monte Carlo metaheuristics is Evolutionary Computation⁹. It encompasses all algorithms that are based on a set of multiple solution candidates (called population) which are iteratively refined. This field of optimization is also a class of Soft Computing¹⁰ as well as a part of the artificial intelligence¹¹ area. Its most important members are evolutionary algorithms and Swarm Intelligence, which will be discussed in-depth in this book.

Besides the evolutionary approaches which are mostly nature-inspired, there exist also methods that copy physical processes like the before-mentioned Simulated Annealing and Parallel Tempering, as well as purely artificial techniques like Tabu Search and Random Optimization.

As a preview of what can be found in this book, we have marked the techniques that will be discussed with a thicker border in Figure 1.1.

1.1.2 Classification According to Properties

The taxonomy just introduced classifies the optimization methods according to their algorithmic structure and underlying principles, in other words, from the viewpoint of theory. A software engineer or a user who wants to solve a problem with such an approach is however more interested in its “interfacing features” such as speed and precision.

Speed and precision are conflicting objectives, at least in terms of probabilistic algorithms. A general rule of thumb is that you can gain improvements in accuracy of optimization only by investing more time. Scientists in the area of global optimization try to push this Pareto frontier¹² further by inventing new approaches and enhancing or tweaking existing ones.

Optimization Speed

When it comes to time constraints and hence, the required speed of the optimization algorithm, we can distinguish two main types of optimization use cases.

Definition 1.4 (Online Optimization). Online optimization problems are tasks that need to be solved quickly in a time span between ten milliseconds to a few minutes. In order to find a solution in this short time, optimality is normally traded in for speed gains.

Examples for online optimization are robot localization, load balancing, services composition for business processes (see for example Section 21.2.1 on page 328), or updating a factory’s machine job schedule after new orders came in. From the examples, it becomes clear that online optimization tasks are often carried out repetitively – new orders will, for instance, continuously arrive in a production facility and need to be scheduled to machines in a way that minimizes the waiting time of all jobs.

Definition 1.5 (Offline Optimization). In offline optimization problems, time is not so important and a user is willing to wait maybe even days if she can get an optimal or close-to-optimal result.

⁹ http://en.wikipedia.org/wiki/Evolutionary_computation [accessed 2007-09-17]

¹⁰ http://en.wikipedia.org/wiki/Soft_computing [accessed 2007-09-17]

¹¹ http://en.wikipedia.org/wiki/Artificial_intelligence [accessed 2007-09-17]

¹² Pareto frontiers will be discussed in Section 1.2.4 on page 13.

Such problems regard for example design optimization, data mining (see for instance Section 21.1 on page 317), or creating long-term schedules for transportation crews. These optimization processes will usually be carried out only once in a long time.

Before doing anything else, one must be sure about to which of these two classes the problem to be solved belongs.

TODO

Number of Criteria

Optimization algorithms can be divided in such which try to find the best values of single objective functions f and such that optimize sets F of target functions. This distinction between single-objective optimization and multi-objective optimization is discussed in depth in Section 1.2.2.

1.2 What is an optimum?

We have already said that global optimization is about finding the best possible solutions for given problems. Thus, it cannot be a bad idea to start out by discussing what it is that makes a solution *optimal*¹³.

1.2.1 Single Objective Functions

In the case of optimizing a single criterion f , an optimum is either its maximum or minimum, depending on what we are looking for. If we own a manufacturing plant and have to assign incoming orders to machines, we will do this in a way that *minimizes* the time needed to complete them. On the other hand, we will arrange the purchase of raw material, the employment of staff, and the placing of commercials in a way that *maximizes* our profit. In global optimization, it is a convention that optimization problems are most often defined as *minimizations* and if a criterion f is subject to maximization, we simply minimize its negation $(-f)$.

Figure 1.2 illustrates such a function f defined over a two-dimensional space $\mathbb{X} = (X_1, X_2)$. As outlined in this graphic, we distinguish between local and global optima. A global optimum is an optimum of the whole domain $\mathbb{X}$ while a local optimum is an optimum of only a subset of $\mathbb{X}$.

Definition 1.6 (Local Maximum). A (local) maximum $\hat{x}_l \in \mathbb{X}$ of one (objective) function $f : \mathbb{X} \mapsto \mathbb{R}$ is an input element with $f(\hat{x}_l) \geq f(x)$ for all x neighboring $\hat{x}_l$.

If $\mathbb{X} \subseteq \mathbb{R}^n$, we can write:

$$\forall \hat{x}_l \exists \varepsilon > 0 : f(\hat{x}_l) \geq f(x) \quad \forall x \in \mathbb{X}, |x - \hat{x}_l| < \varepsilon \quad (1.1)$$

Definition 1.7 (Local Minimum). A (local) minimum $\check{x}_l \in \mathbb{X}$ of one (objective) function $f : \mathbb{X} \mapsto \mathbb{R}$ is an input element with $f(\check{x}_l) \leq f(x)$ for all x neighboring $\check{x}_l$.

If $\mathbb{X} \subseteq \mathbb{R}$, we can write:

$$\forall \check{x}_l \exists \varepsilon > 0 : f(\check{x}_l) \leq f(x) \quad \forall x \in \mathbb{X}, |x - \check{x}_l| < \varepsilon \quad (1.2)$$

Definition 1.8 (Local Optimum). A (local) optimum $x_l^* \in \mathbb{X}$ of one (objective) function $f : \mathbb{X} \mapsto \mathbb{R}$ is either a local maximum or a local minimum.

¹³ http://en.wikipedia.org/wiki/Maxima_and_minima [accessed 2007-07-03]




Fig. 1.2: Global and local optima of a two-dimensional function.

Definition 1.9 (Global Maximum). A global maximum $\hat{\mathbf{x}} \in \mathbb{X}$ of one (objective) function $f : \mathbb{X} \mapsto \mathbb{R}$ is an input element with $f(\hat{\mathbf{x}}) \geq f(x) \forall x \in \mathbb{X}$.

Definition 1.10 (Global Minimum). A global minimum $\tilde{\mathbf{x}} \in \mathbb{X}$ of one (objective) function $f : \mathbb{X} \mapsto \mathbb{R}$ is an input element with $f(\tilde{\mathbf{x}}) \leq f(x) \forall x \in \mathbb{X}$.

Definition 1.11 (Global Optimum). A global optimum $\mathbf{x}^* \in \mathbb{X}$ of one (objective) function $f : \mathbb{X} \mapsto \mathbb{R}$ is either a global maximum or a global minimum.

Even a one-dimensional function $f : \mathbb{X} = \mathbb{R} \mapsto \mathbb{R}$ may have more than one global maximum, multiple global minima, or even both in its domain $\mathbb{X}$. Take the cosine function for example: It has global maxima $\hat{\mathbf{x}}_i$ at $\hat{\mathbf{x}}_i = 2i\pi$ and global minima $\tilde{\mathbf{x}}_i$ at $\tilde{\mathbf{x}}_i = (2i+1)\pi$ for all $i \in \mathbb{Z}$. The correct solution of such an optimization problem would then be a set $\mathbf{X}^*$ of all optimal inputs in $\mathbb{X}$ rather than a single maximum or minimum. Furthermore, the exact meaning of *optimal* is problem dependent. In single-objective optimization, it either means minimum or maximum. In multi-objective optimization, there exist a variety of approaches to define optima which we will discuss in-depth in Section 1.2.2.

Definition 1.12 (Optimal Set). The optimal set $\mathbf{X}^*$ is the set that contains all optimal elements.

There are normally multiple, often even infinite many optimal solutions. Since the memory of our computers is limited, we can find only a finite (sub-)set of them. We thus distinguish between the global optimal set $\mathbf{X}^*$ and the set X^* of (seemingly optimal) elements which an optimizer returns. The tasks of global optimization algorithms are

1. to find solutions that are as good as possible and
2. that are also widely different from each other [12].

The second goal becomes obvious if we assume that we have an objective function $f : \mathbb{R} \mapsto \mathbb{R}$ which is optimal for all $x \in [0, 10] \Leftrightarrow x \in \mathbf{X}^*$. This interval contains uncountable many solutions, and an optimization algorithm may yield $X_1^* = \{0, 0.1, 0.11, 0.05, 0.01\}$ or $X_2^* = \{0, 2.5, 5, 7.5, 10\}$ as result. Both sets only represent a small subset of the possible solutions. The second result (X_2^*), however, gives us a broader view on the optimal set.

Even good optimization algorithms do not necessarily find the real global optima but may only be able to approximate them. In other words, $X_3^* = \{-0.3, 5, 7.5, 11\}$ is also a possible result of the optimization process, although containing two sub-optimal elements.

In Chapter 18 on page 263, we will introduce different algorithms and approaches that can be used to maintain an optimal set or to select the optimal elements from a given set during an optimization process.

1.2.2 Multiple Objective Functions

Global optimization techniques are not just used for finding the maxima or minima of single functions f . In many real-world design or decision making problems, they are rather applied to sets F consisting of $n = |F|$ objective functions f_i , each representing one criterion to be optimized [13, 14, 15].

$$F = \{f_i : \mathbb{X} \mapsto Y_i : 0 < i \leq n, Y_i \subseteq \mathbb{R}\} \quad (1.3)$$

Algorithms designed to optimize such sets of objective functions are usually named with the prefix *multi-objective*, like multi-objective evolutionary algorithms which are discussed in Definition 2.2 on page 68.

Factory Example

Multi-objective optimization often means to compromise conflicting goals. If we go back to our factory example, we can specify the following objectives that all are subject to optimization:

- Minimize the time between an incoming order and the shipment of the corresponding product.
- Maximize profit.
- Minimize costs for advertising, personal, raw materials etc..
- Maximize product quality.
- Minimize negative impact on environment.

The last two objectives seem to contradict clearly the cost minimization. Between the personal costs and the time needed for production and the product quality there should also be some kind of (contradictive) relation. The exact mutual influences between objectives can apparently become complicated and are not always obvious.

Artificial Ant Example

Another example for such a situation is the Artificial Ant problem¹⁴ where the goal is to find the most efficient controller for a simulated ant. The efficiency of an ant should not only be measured by the amount of food it is able to pile. For every food item, the ant needs to walk to some point. The more food it piles, the longer the distance it needs to walk. If its behavior is driven by a clever program, it may walk along a shorter route which would not be discovered by an ant with a clumsy controller. Thus, the distance it has to cover to find the food or the time it needs to do so may also be considered in the optimization process. If two control programs produce the same results and one is smaller (i. e., contains fewer instructions) than the other, the smaller one should be preferred. Like in the factory example, the optimization goals conflict with each other.

From these both examples, we can gain another insight: To find the global optimum could mean to maximize one function $f_i \in F$ and to minimize another one $f_j \in F$, ($i \neq j$). Hence, it makes no sense to talk about a *global maximum* or a *global minimum* in terms

¹⁴ See Section 20.3.1 on page 303 for more details.

of multi-objective optimization. We will thus retreat to the notation of the set of optimal elements $\mathbf{x}^* \in X^* \subseteq \mathbb{X}$.

Since compromises for conflicting criteria can be defined in many ways, there exist multiple approaches to define what an optimum is. These different definitions, in turn, lead to different sets $\mathbf{X}^*$.

Graphical Example 1




Fig. 1.3: Two functions f_1 and f_2 with different maxima $\hat{\mathbf{x}}_1$ and $\hat{\mathbf{x}}_2$.

We will discuss some of these approaches in the following by using two graphical examples for illustration purposes. In the first example pictured in Figure 1.3, we want to maximize two independent objective functions $F_1 = \{f_1, f_2\}$. Both objective functions have the real numbers $\mathbb{R}$ as problem space $\mathbb{X}_1$. The maximum (and thus, the optimum) of f_1 is $\hat{\mathbf{x}}_1$ and the largest value of f_2 is at $\hat{\mathbf{x}}_2$. In Figure 1.3, we can easily see that f_1 and f_2 are partly conflicting: Their maxima are at different locations and there even exist areas where f_1 rises while f_2 falls and vice versa.

Graphical Example 2




Fig. 1.4: Two functions f_3 and f_4 with different minima $\check{\mathbf{x}}_1$, $\check{\mathbf{x}}_2$, $\check{\mathbf{x}}_3$, and $\check{\mathbf{x}}_4$.

The objective functions f_1 and f_2 in the first example are mappings of a one-dimensional problem space $\mathbb{X}_1$ to the real numbers that are to be maximized. In the second exam-

ple sketched in Figure 1.4, we instead minimize two functions f_3 and f_4 that map a two-dimensional problem space $\mathbb{X}_2 \subset \mathbb{R}^2$ to the real numbers $\mathbb{R}$. Both functions have two global minima; the lowest values of f_3 are $\tilde{\mathbf{x}}_1$ and $\tilde{\mathbf{x}}_2$ whereas f_4 gets minimal at $\tilde{\mathbf{x}}_3$ and $\tilde{\mathbf{x}}_4$. It should be noted that $\tilde{\mathbf{x}}_1 \neq \tilde{\mathbf{x}}_2 \neq \tilde{\mathbf{x}}_3 \neq \tilde{\mathbf{x}}_4$.

1.2.3 Weighted Sum

The simplest method to define what is optimal is computing a weighted sum $g(x)$ of all the functions $f_i(x) \in F$. Each objective f_i is multiplied with a weight w_i representing its importance. Using signed weights also allows us to minimize one objective and to maximize another. We can, for instance, apply a weight $w_a = 1$ to an objective function f_a and the weight $w_b = -1$ to the criterion f_b . By minimizing $g(x)$, we then actually minimize the first and maximize the second objective function. If we instead maximize $g(x)$, the effect would be converse and f_b would be minimized and f_a would be maximized.

Either way, multi-objective problems are reduced to single-objective ones by this method.

$$g(x) = \sum_{i=1}^n w_i f_i(x) = \sum_{\forall f_i \in F} w_i f_i(x) \quad (1.4)$$

$$\mathbf{x}^* \in \mathbf{X}^* \Leftrightarrow g(\mathbf{x}^*) \geq g(x) \quad \forall x \in \mathbb{X} \quad (1.5)$$

Graphical Example 1

Figure 1.5 demonstrates optimization with the weighted sum approach for the example given in Section 1.2.2. The weights are both set to 1 = $w_1 = w_2$. If we maximize $g_1(2)$, we will thus also maximize the functions f_1 and f_2 . This leads to a single optimum $\mathbf{x}^* = \hat{\mathbf{x}}$.




Fig. 1.5: Optimization using the weighted sum approach (first example).

Graphical Example 2

The sum of the two-dimensional functions f_3 and f_4 from the second graphical example given in Section 1.2.2 is sketched in Figure 1.6. Again we set the weights w_3 and w_4 to 1. The sum g_2 however is subject to minimization. The graph of g_2 has two especially deep valleys. At the bottoms of these valleys, the two global minima $\tilde{\mathbf{x}}_5$ and $\tilde{\mathbf{x}}_6$ can be found.




Fig. 1.6: Optimization using the weighted sum approach (second example).

Problems with Weighted Sums

The drawback of this approach is that it cannot handle functions that rise or fall with different speed¹⁵ properly. In Figure 1.7, we have sketched the sum $g(x)$ of the two objective functions $f_1(x) = -x^2$ and $f_2(x) = e^{x-2}$. If we minimize or maximize this sum, we will always




Fig. 1.7: A problematic constellation for the weighted sum approach.

disregard one of two functions, depending on the interval chosen. For small x , f_2 is negligible compared to f_1 . For $x > 5$ it begins to outpace f_1 which, in turn, will now become negligible. Such functions cannot be added up properly using constant weights. Even if we would set w_1 to the really large number 10^{10} , f_1 will become insignificant for all $x > 40$, because

¹⁵ See Section 29.1.3 on page 482

or http://en.wikipedia.org/wiki/Asymptotic_notation [accessed 2007-07-03] for related information.

$\left| \frac{-(40^2) \cdot 10^{10}}{e^{40-2}} \right| \approx 0.0005$. Therefore, weighted sums are only suitable to optimize functions that at least share the same big-**O** notation (see Section 29.1.3 on page 482). Often, it is not obvious to us how objective functions will fall or rise. How can we, for instance, determine whether the objective maximizing the food piled by an Artificial Ant rises in comparison to the objective minimizing the distance walked by the simulated insect?

1.2.4 Pareto Optimization

Pareto efficiency¹⁶ (also called Pareto optimality) is an important notion in neoclassical economics with broad applications in game theory, engineering and the social sciences [16, 17]. It defines the frontier of solutions that can be reached by trading-off conflicting objectives in an optimal manner. From this front, a decision maker (be it a human or another algorithm) can finally choose the configurations that, in his opinion, suite best [18, 19, 20, 21, 22]. The notation of optimal in the sense of Pareto efficiency is strongly based on the definition of domination:

Definition 1.13 (Domination). An element x_1 dominates (is preferred to) an element x_2 ($x_1 \vdash x_2$) if x_1 is better than x_2 in at least one objective function and not worse with respect to all other objectives. Based on the set F of objective functions f , we can write:

$$x_1 \vdash x_2 \Leftrightarrow \forall i : 0 < i \leq n \Rightarrow \omega_i f_i(x_1) \leq \omega_i f_i(x_2) \wedge \exists j : 0 < j \leq n : \omega_j f_j(x_1) < \omega_j f_j(x_2) \quad (1.6)$$

$$\omega_i = \begin{cases} 1 & \text{if } f_i \text{ should be minimized} \\ -1 & \text{if } f_i \text{ should be maximized} \end{cases} \quad (1.7)$$

Different from the weights in the weighted sum approach, the factors ω_i only carry sign information which allows us to maximize some objectives and to minimize some other criteria.

The Pareto domination relation defines a strict partial order (see Definition 26.31 on page 415) on the space of possible objective values. In contrast, the weighted sum approach imposes a total order by projecting it into the real numbers $\mathbb{R}$.

Definition 1.14 (Pareto Optimal). An element $\mathbf{x}^* \in \mathbb{X}$ is Pareto optimal (and hence, part of the optimal set $\mathbf{X}^*$) if it is not dominated by any other element in the problem space $\mathbb{X}$. In terms of Pareto optimization, $\mathbf{X}^*$ is called the Pareto set or the Pareto Frontier.

$$\mathbf{x}^* \in \mathbf{X}^* \Leftrightarrow \nexists x \in \mathbb{X} : x \vdash \mathbf{x}^* \quad (1.8)$$

Graphical Example 1

In Figure 1.8, we illustrate the impact of the definition of Pareto optimality on our first example (outlined in Section 1.2.2). We assume again that f_1 and f_2 should both be maximized and hence, $\omega_1 = \omega_2 = -1$. The areas shaded with dark gray are Pareto optimal and thus, represent the optimal set $\mathbf{X}^* = [x_2, x_3] \cup [x_5, x_6]$ which here contains infinite many elements¹⁷. All other points are dominated, i. e., not optimal.

The points in the area between x_1 and x_2 (shaded in light gray) are dominated by other points in the same region or in $[x_2, x_3]$, since both functions f_1 and f_2 can be improved by increasing x . If we start at the leftmost point in $\mathbb{X}$ (which is position x_1), for instance, we can go one small step Δ to the right and will find a point $x_1 + \Delta$ dominating x_1 because $f_1(x_1 + \Delta) > f_1(x_1)$ and $f_2(x_1 + \Delta) > f_2(x_1)$. We can repeat this procedure and will always

¹⁶ http://en.wikipedia.org/wiki/Pareto_efficiency [accessed 2007-07-03]

¹⁷ In practice, of course, our computers can only handle finitely many elements




Fig. 1.8: Optimization using the Pareto Frontier approach.

find a new dominating point until we reach x_2 . x_2 demarks the global maximum of f_2 , the point with the highest possible f_2 value, which cannot be dominated by any other point in $\mathbb{X}$ by definition (see Equation 1.6).

From here on, f_2 will decrease for a while, but f_1 keeps rising. If we now go a small step Δ to the right, we will find a point $x_2 + \Delta$ with $f_2(x_2 + \Delta) < f_2(x_2)$ but also $f_1(x_2 + \Delta) > f_1(x_2)$. One objective can only get better if another one degenerates. In order to increase f_1 , f_2 would be decreased and vice versa and so the new point is not dominated by x_2 . Although some of the $f_2(x)$ values of the other points $x \in [x_1, x_2]$ may be larger than $f_2(x_2 + \Delta)$, $f_1(x_2 + \Delta) > f_1(x)$ holds for all of them. This means that no point in $[x_1, x_2]$ can dominate any point in $[x_2, x_4]$ because f_1 keeps rising until x_4 is reached.

At x_3 however, f_2 steeply falls to a very low level. A level lower than $f_2(x_5)$. Since the f_1 values of the points in $[x_5, x_6]$ are also higher than those of the points in $(x_3, x_4]$, all points in the set $[x_5, x_6]$ (which also contains the global maximum of f_1) dominate those in $(x_3, x_4]$. For all the points in the white area between x_4 and x_5 and after x_6 , we can derive similar relations. All of them are also dominated by the non-dominated regions that we have just discussed.

Graphical Example 2

Another method to visualize the Pareto relationship is outlined in Figure 1.9 for our second graphical example. For a certain resolution of the problem space $\mathbb{X}_2$, we have counted the number of elements that dominate each element $x \in \mathbb{X}_2$. The higher this number, the worst is the element x in terms of Pareto optimization. Hence, those solution candidates residing in the valleys of Figure 1.9 are better than those which are part of the hills. This Pareto ranking approach is also used in many optimization algorithms as part of the fitness assignment scheme (see Section 2.3.3 on page 83, for instance). A non-dominated element is, as the name says, not dominated by any other solution candidate. These elements are Pareto optimal and have a domination-count of zero. In Figure 1.9, there are four such areas $\mathbf{X}_1^*$, $\mathbf{X}_2^*$, $\mathbf{X}_3^*$, and $\mathbf{X}_4^*$.

If we compare Figure 1.9 with the plots of the two functions f_3 and f_4 in Figure 1.4, we can see that hills in the domination space occur at positions where both, f_3 and f_4 have high values. Conversely, regions of the problem space where both functions have small values are dominated by very few elements.

Besides these examples here, another illustration of the domination relation which may help understanding Pareto optimization can be found in Section 2.3.3 on page 83 (Figure 2.4 and Table 2.1).




Fig. 1.9: Optimization using the Pareto Frontier approach (second example).

Problems of Pure Pareto Optimization

The complete Pareto optimal set is often not the wanted result of an optimization algorithm. Usually, we are rather interested in some special areas of the Pareto front only.

Artificial Ant Example

We can again take the Artificial Ant example to visualize this problem. In Section 1.2.2 on page 9 we have introduced multiple conflicting criteria in this problem.

- Maximize the amount of food piled.
- Minimize the distance covered or the time needed to find the food.
- Minimize the size of the program driving the ant.

Pareto optimization may now yield for example:

- A program consisting of 100 instructions, allowing the ant to gather 50 food items when walking a distance of 500 length units.
- A program consisting of 100 instructions, allowing the ant to gather 60 food items when walking a distance of 5000 length units.
- A program consisting of 10 instructions, allowing the ant to gather 1 food item when walking a distance of 5 length units.
- A program consisting of 0 instructions, allowing the ant to gather 0 food item when walking a distance of 0 length units.

The result of the optimization process obviously contains two useless but non-dominated individuals which occupy space in the population and the non-dominated set. We also invest processing time in evaluating them, and even worse, they may dominate solutions that are not optimal but fall into the space behind the interesting part of the Pareto front. Furthermore, memory restrictions usually force us to limit the size of the list of non-dominated solutions found during the search. When this size limit is reached, some optimization algorithms use a clustering technique to prune the optimal set while maintaining diversity. On one hand, this is good since it will preserve a broad scan of the Pareto frontier. In this case on the other hand, a short but dumb program is of course very different from a longer, intelligent one. Therefore, it will be kept in the list and other solutions which differ less from each other but are more interesting for us will be discarded.

Furthermore, non-dominated elements have a higher probability of being explored further. This then leads inevitably to the creation of a great proportion of useless offspring. In the next generation, these useless offspring will need a good share of the processing time to be evaluated.

Thus, there are several reasons to force the optimization process into a wanted direction. In Section 21.2.2 on page 334 you can find an illustrative discussion on the drawbacks of strict Pareto optimization in a practical example (evolving web service compositions).

1.2.5 The Method of Inequalities

One approach for dealing with these problems is the *Method of Inequalities* (MOI) [23, 24, 25] which has its roots in operations research.

This approach is applied by explicitly specifying an area of interest in form of a goal range $[\check{r}_i, \hat{r}_i]$ for each objective function f_i . Based on inequalities we can define three categories of solution candidates and each element $x \in \mathbb{X}$ belongs to one of them:

1. It fulfills all of the goals, i. e.,

$$\check{r}_i \leq f_i(x) \leq \hat{r}_i \quad \forall i \in [1, |F|] \quad (1.9)$$

2. It fulfills some (but not all) of the goals, i. e.,

$$(\exists i \in [1, |F|] : \check{r}_i \leq f_i(x) \leq \hat{r}_i) \wedge (\exists j \in [1, |F|] : (f_j(x) < \check{r}_j) \vee (f_j(x) > \hat{r}_j)) \quad (1.10)$$

3. It fulfills none of the goals, i. e.,

$$(f_i(x) < \check{r}_i) \vee (f_i(x) > \hat{r}_i) \quad \forall i \in [1, |F|] \quad (1.11)$$

Using these groups we can create a new comparison mechanism:

- The solution candidates that fulfill all goals are preferred instead of all other individuals that either fulfill some or no goals.
- The solution candidates that are not able to fulfill any of the goals succumb to those which fulfill at least some goals.
- Only the solutions that are in the same group are compared on basis on the Pareto domination relation.

By doing so, the optimization process will be driven into the direction of the interesting part of the Pareto frontier. Less effort will be spent in creating and evaluating individuals in parts of the problem space that most probably do not contain any valid solution.

Graphical Example 1

In Figure 1.10, we apply the Method of Inequalities to our first graphical example. We impose the same goal ranges on both objectives $\hat{r}_1 = \hat{r}_2$ and $\check{r}_1 = \check{r}_2$. By doing so, the second non-dominated region from the Pareto example Figure 1.8 suddenly becomes infeasible, since f_1 rises over $\hat{r}_1$ there. Also, the greater part of the first optimal area from this example is infeasible because f_2 drops under $\check{r}_2$. In the whole domain $\mathbb{X}$ of the optimization problem, only the regions $[x_1, x_2]$ and $[x_3, x_4]$ fulfill all the target criteria. To these elements, Pareto comparisons are applied. It turns out that the elements in $[x_3, x_4]$ dominate all the elements $[x_1, x_2]$ since they provide higher values in f_1 for same values in f_2 . If we scan through $[x_3, x_4]$ from left to right, we can see the f_1 rises while f_2 degenerates, which is why the elements in this area cannot be dominated each other and, hence, are all optimal.

Graphical Example 2

In Figure 1.11 we apply the Method of Inequalities to our second graphical example from Section 1.2.2. We apply two different ranges of interest $[\check{r}_3, \hat{r}_3]$ and $[\check{r}_4, \hat{r}_4]$ on f_3 and f_4 as sketched in Fig. 1.11.1.




Fig. 1.10: Optimization using the Method of Inequalities approach (first example).




Fig. 1.11.1: The ranges applied to f_3 and f_4 .




Fig. 1.11.2: The Method of Inequalities class division.




Fig. 1.11.3: The Method of Inequalities ranking.

Fig. 1.11: Optimization using the Method of Inequalities approach (first example).

Like we did in the second example for Pareto optimization, we want to plot the quality of the elements in the problem space. Therefore, we first assign a number $c \in \{1, 2, 3\}$ to each of its elements in Fig. 1.11.2. This number corresponds to the classes to which the elements belong, i.e., 1 means that a solution candidate fulfills all inequalities, for an element of class 2, at least some of the constraints hold, and the elements in class 3 fail all requirements. Based on this class division, we can then perform a modified Pareto counting where each element dominates all the elements in higher classes Fig. 1.11.3. The result is that multiple single optima $\mathbf{x}_1^*$, $\mathbf{x}_2^*$, $\mathbf{x}_3^*$, etc., and even a set of adjacent, non-dominated elements $\mathbf{X}_g^*$ occurs. These elements are, again, situated at the bottom of the illustrated landscape whereas the worst solution candidates reside on hill tops.

Limitations and Other Methodes

If we want to apply the Method of Inequalities, we need to know the regions of interest in the objective space beforehand. While this is the case in most applications, a real drawback is that we must confine the possible results of the optimization process to one single rectangular window.

Other techniques similar to the Method of Inequalities are *Goal Attainment* [26] and *Goal Programming*¹⁸ [27]. Often, characteristics of these methods are added to evolutionary algorithms [28, 29, 30, 31, 32].

TODO

1.2.6 External Decision Maker

In order to circumvent the limitations of the previous approaches, Fonseca and Fleming [33; 15] introduced their more general concept of an External Decision Maker which (or who) decides which solution candidates prevail. The basic idea behind this is that Pareto optimization imposes only a partial order¹⁹ on the solution candidates. In a partial order, elements may exist which neither succeed nor precede each other. As we have seen in Section 1.2.4, there can, for instance, be two individuals $x_1, x_2 \in \mathbb{X}$ with neither $x_1 \vdash x_2$ nor $x_2 \vdash x_1$. A special case of this situation is the non-dominated set, the so-called *Pareto frontier* which we try to estimate with the optimization process.

Most fitness assignment processes, however, require some sort of total order²⁰, where each individual is either better or worse than each other (except for the case of identical solution candidates which are, of course, equal to each other). The fitness assignment algorithms can create such a total order by themselves. One example for doing this is the Pareto ranking which we will discuss later in Section 2.3.3 on page 83, where the number of individuals dominating a solution candidate denotes its fitness.

While this method of ordering is a good default approach able of directing the search into the direction of the Pareto frontier and delivering a broad scan of it, it neglects the fact that the user of the optimization most often is not interested in the whole optimal set but has *preferences*, certain regions of interest [34]. This region will then exclude the infeasible (but Pareto optimal) programs for the Artificial Ant as discussed in Section 1.2.4. What the user wants is a detailed scan of these areas, which often cannot be delivered by pure Pareto optimization.

Here comes the External Decision Maker as an expression of the user's preferences [35] into play, as illustrated in Figure 1.12. The task of this decision maker is to provide a cost function $u : \mathbb{Y} \mapsto \mathbb{R}$ (or utility function, if the underlying optimizer is maximizing) which

¹⁸ http://en.wikipedia.org/wiki/Goal_programming [accessed 2007-07-03]

¹⁹ A definition of *partial order* relations is specified in Definition 26.31 on page 415.

²⁰ The concept of *total orders* is elucidated in Definition 26.32 on page 415.




Fig. 1.12: An external decision maker providing an evolutionary algorithm with utility values.

maps the space of objective values $\mathbb{Y}$ (which is usually $\mathbb{R}^n$) to the space of real numbers $\mathbb{R}$. Since there is a total order defined on the real numbers, this process is another way of resolving the “incomparability-situation”. The structure of the decision making process u can freely be defined and may incorporate any of the previously mentioned methods. u could, for example, be reduced to compute a weighted sum of the objective values, to perform an implicit Pareto ranking, or to compare individuals based on pre-specified goal-vectors. Furthermore, it may even incorporate forms of artificial intelligence, other forms of multi-criterion Decision Making, and even interaction with the user. This technique allows focusing the search onto solutions which are not only optimal in the Pareto sense, but also feasible and interesting from the viewpoint of the user.

Fonseca and Fleming make a clear distinction between fitness and cost values. Cost values have some meaning outside the optimization process and are based on user preferences. Fitness values on the other hand are an internal construct of the search with no meaning outside the optimizer (see Definition 1.35 on page 28 for more details). If External Decision Makers are applied in evolutionary algorithms or other search paradigms that are based on fitness measures, these will be computed using the values of the cost function instead of the objective functions [36, 35, 37].

1.2.7 Prevalence Optimization

We have now discussed various approaches which define optima in terms of multi-objective optimization and steer the search process into their direction. Let us subsume all of them in general approach. From the concept of Pareto optimization to the Method of Inequalities, the need to *compare* elements of the problem space in terms of their quality as solution for a given problem winds like a red thread through this matter. Even the weighted sum approach and the External Decision Maker do nothing else than mapping multi-dimensional vectors to the real numbers in order to make them comparable.

If we compare two solution candidates x_1 and x_2 , either x_1 is better than x_2 , vice versa, or both are of equal quality. Hence, there are three possible relations between two elements of the problem space. These two results can be expressed with a comparator function cmp_F .

Definition 1.15 (Comparator Function). A comparator function $\text{cmp} : \mathbb{A}^2 \mapsto \mathbb{R}$ maps all pairs of elements $(a_1, a_2) \in \mathbb{A}^2$ to the real numbers $\mathbb{R}$ according to two complementing partial orders²¹ R_1 and R_2 :

$$R_1(a_1, a_2) \Leftrightarrow \text{cmp}(a_1, a_2) < 0 \quad \forall a_1, a_2 \in \mathbb{A} \quad (1.12)$$

$$R_2(a_1, a_2) \Leftrightarrow \text{cmp}(a_1, a_2) > 0 \quad \forall a_1, a_2 \in \mathbb{A} \quad (1.13)$$

$$\overline{R_1(a_1, a_2)} \wedge \overline{R_2(a_1, a_2)} \Leftrightarrow \text{cmp}(a_1, a_2) = 0 \quad \forall a_1, a_2 \in \mathbb{A} \quad (1.14)$$

$$\text{cmp}(a, a) = 0 \quad \forall a \in \mathbb{A} \quad (1.15)$$

²¹ Partial orders are introduced in Definition 26.30 on page 415.

R_1 (and hence, $\text{cmp}(a_1, a_2) < 0$) is equivalent to the precedence relation and R_2 denotes succession.

From the three defining equations, many features of cmp can be deduced. It is, for instance, transitive, i. e., $\text{cmp}(a_1, a_2) < 0 \wedge \text{cmp}(a_2, a_3) < 0 \Rightarrow \text{cmp}(a_1, a_3) < 0$. Provided with the knowledge of the objective functions $f \in F$, such a comparator function cmp_F can be imposed on the problem spaces of our optimization problems:

Definition 1.16 (Prevalence Comparator Function). A prevalence comparator function $\text{cmp}_F : \mathbb{X}^2 \mapsto \mathbb{R}$ maps all pairs $(x_1, x_2) \in \mathbb{X}^2$ of solution candidates to the real numbers $\mathbb{R}$ according to Definition 1.15.

The subscript F in cmp_F illustrates that the comparator has access to all the values of the objective functions in addition to the problem space elements which are its parameters. As shortcut for this comparator function, we introduce the prevalence notation as follows:

Definition 1.17 (Prevalence). An element x_1 prevails over an element x_2 ($x_1 \succ x_2$) if the application-dependent prevalence comparator function $\text{cmp}_F(x_1, x_2) \in \mathbb{R}$ returns a value less than 0.

$$(x_1 \succ x_2) \Leftrightarrow \text{cmp}_F(x_1, x_2) < 0 \quad \forall x_1, x_2 \in \mathbb{X} \quad (1.16)$$

$$(x_1 \succ x_2) \wedge (x_2 \succ x_3) \Rightarrow x_1 \succ x_3 \quad \forall x_1, x_2, x_3 \in \mathbb{X} \quad (1.17)$$

It is easy to see that we can define Pareto domination relations and Method of Inequalities-based comparisons, as well as the weighted sum combination of objective values based on this notation. Together with the fitness assignment strategies which will be introduced later in this book (see Section 2.3 on page 82), it covers many of the most sophisticated multi-objective techniques that are proposed, for instance, in [33, 38, 28]. By replacing the Pareto approach with prevalence comparisons, all the optimization algorithms (especially many of the evolutionary techniques) relying on domination relations can be used in their original form while offering the new ability of scanning special regions of interests of the optimal frontier.

Since the comparator function cmp_F and the prevalence relation impose a partial order on the problem space $\mathbb{X}$ like the domination relation does, we can construct the optimal set in a way very similar to Equation 1.8:

$$\mathbf{x}^* \in \mathbf{X}^* \Leftrightarrow \nexists x \in \mathbb{X} : x \neq \mathbf{x}^* \wedge x \succ \mathbf{x}^* \quad (1.18)$$

For illustration purposes, we will exercise the prevalence approach on the examples of the weighted sum $\text{cmp}_{F,F,\text{weightedS}}$ method²² with the weights w_i as well as on the domination-based Pareto optimization²³ $\text{cmp}_{F,\text{Pareto}}$ with the objective directions ω_i :

$$\text{cmp}_{F,\text{weightedS}}(x_1, x_2) = \sum_{i=1}^{|F|} (w_i f_i(x_2) - w_i f_i(x_1)) \equiv g(x_2) - g(x_1) \quad (1.19)$$

$$\text{cmp}_{F,\text{Pareto}}(x_1, x_2) = \begin{cases} -1 & \text{if } x_1 \vdash x_2 \\ 1 & \text{if } x_2 \vdash x_1 \\ 0 & \text{otherwise} \end{cases} \quad (1.20)$$

²² See Equation 1.4 on page 11 for more information on weighted sum optimization.

²³ Pareto optimization was defined in Equation 1.6 on page 13.

Artificial Ant Example

With the prevalence comparator, we can also easily solve the problem stated in Section 1.2.4 by no longer encouraging the evolution of useless programs for Artificial Ants while retaining the benefits of Pareto optimization. The comparator function simple can be defined in a way that they will always be prevailed by useful programs. It therefore may incorporate the knowledge on the importance of the objective functions. Let f_1 be the objective function with an output proportional to the food piled, f_2 would denote the distance covered in order to find the food, and f_3 would be the program length. Equation 1.21 demonstrates one possible comparator function for the Artificial Ant problem.

$$\text{cmp}_{F,ant}(x_1, x_2) = \begin{cases} -1 & \text{if } (f_1(x_1) > 0 \wedge f_1(x_2) = 0) \vee \\ & (f_2(x_1) > 0 \wedge f_2(x_2) = 0) \vee \\ & (f_3(x_1) > 0 \wedge f_1(x_2) = 0) \\ 1 & \text{if } (f_1(x_2) > 0 \wedge f_1(x_1) = 0) \vee \\ & (f_2(x_2) > 0 \wedge f_2(x_1) = 0) \vee \\ & (f_3(x_2) > 0 \wedge f_1(x_1) = 0) \\ \text{cmp}_{F,Pareto}(x_1, x_2) & \text{otherwise} \end{cases} \quad (1.21)$$

Later in this book, we will discuss some of the most popular optimization strategies. Although they are usually implemented based on Pareto optimization, we will always introduce them using prevalence.

1.3 The Structure of Optimization

After we have discussed what optima are and have seen a crude classification of global optimization algorithms, let us now take a look on the general structure common to all optimization processes. This structure consists of a number of well defined spaces and sets and the mappings between them. Based on this structure of optimization, we will introduce the abstractions *fitness landscapes*, *problem landscape*, and *optimization problem* which will lead us to a more thorough definition of what optimization is.

1.3.1 Spaces, Sets, and Elements

In this section, we elaborate on the relation between the (possibly different) representations of solution candidates for search and for evaluation. We will show how these representations are connected and introduce *fitness* as a relative utility measures defined on sets of solution candidates. You will find that the general model introduced here applies to all the global optimization methods mentioned in this book, often in a simplified manner. One example for the structure of optimization is outlined in Figure 1.13 using a genetic algorithm which binary encodes the points in a plane as illustration.

The Problem Space and the Solutions therein

Whenever we tackle an optimization problem, we first have to define the type of the possible solutions. For deriving a controller for the Artificial Ant problem, we could choose programs or artificial neural networks as solution representation. If we are to find the root of a mathematical function, we would go for real numbers as solution candidates and when configuring and customizing a car for a sales offer, we are looking for an optimal feature selection with costs below a given price. With this initial restriction to a certain type of results, we have specified the problem space $\mathbb{X}$.




Fig. 1.13: Spaces, Sets, and Elements involved in an optimization process.

Definition 1.18 (Problem Space). The problem space $\mathbb{X}$ of a certain optimization problem is the set containing all elements x which could be its solution.

In dependence on genetic algorithms, the problem space is often synonymously referred to as *phenome*. The problem space $\mathbb{X}$ is often restricted by

- *logical constraints* that rule out elements which cannot be solutions, like programs of zero length when trying to solve the Artificial Ant problem and
- *practical constraints* that prevent us, for instance, from taking *all* real numbers into consideration in the minimization process of a real function. On our off-the-shelf CPUs or with the Java programming language, we can only use 64 bit floating point numbers. With these 64 bit, it is only possible to express numbers up to a certain precision and we cannot have more than 15 or so decimals.

Definition 1.19 (Solution Candidate). A solution candidate x is an element of the problem space $\mathbb{X}$ of a certain optimization problem.

We refer to solution candidates synonymously as *phenotypes*. Somewhere inside the problem space, the solutions of the optimization problem will be located (if the problem can actually be solved, that is).

Definition 1.20 (Solution Space). We call the union of all solutions to an optimization problem its solution space $\mathbb{S}$. This solution space contains (and can be equal to) the global optimal set $\mathbf{X}^*$. There may exist valid solutions $x \in \mathbb{S}$ which are not elements of the $\mathbf{X}^*$.

$$\mathbf{X}^* \subseteq \mathbb{S} \subseteq \mathbb{X} \quad (1.22)$$

The Search Space

As previously mentioned, the type of the solution candidates depends on the problem to be solved. Since there are many different applications for optimization, there are many different forms of problem spaces. It would be cumbersome to define search operations time and again for each new problem space we encounter. Such an approach would not only be error-prone, it would also make it very hard to formulate general laws or consolidated findings. It therefore makes sense to distinguish between the search space $\mathbb{G}$ and the problem space $\mathbb{X}$.

Definition 1.21 (Search Space). The search space $\mathbb{G}$ of a certain optimization problem is the set of elements g on which the search operations of an optimization algorithm work.

In dependence on genetic algorithms, we often refer to the search space synonymously as *genome*²⁴, a term coined by the German biologist Winkler [39] as a portmanteau of the words gene and chromosome [40]. The genome is the whole hereditary information of organisms. This includes both, the genes and the non-coding sequences of the Deoxyribonucleic acid (DNA²⁵), which is illustrated in Figure 1.14. Simply put, the DNA is a string of base pairs that encodes the phenotypical characteristics of the creature it belongs to.

Definition 1.22 (Genotype). The elements $g \in \mathbb{G}$ of the search space $\mathbb{G}$ of a given problem are called the genotypes.

The elements of the search space rarely are unstructured aggregations. Instead, they often consist of distinguishable parts, hierarchical units, or well-typed data structures. The same goes for the DNA in biology. It consists of *genes*, segments of nucleic acid, that contain the information necessary to produce RNA strings in a controlled manner²⁶. A fish, for

²⁴ <http://en.wikipedia.org/wiki/Genome> [accessed 2007-07-15]

²⁵ <http://en.wikipedia.org/wiki/Dna> [accessed 2007-07-03]

²⁶ <http://en.wikipedia.org/wiki/Gene> [accessed 2007-07-03]




Fig. 1.14: A sketch of a part of a DNA molecule.

instance, may have a gene for the color of its scales. This gene, in turn, could have two possible “values” called *alleles*²⁷, determining whether the scales will be brown or gray. The genetic algorithm community has adopted this notation long ago and we can use it for arbitrary search spaces.

Definition 1.23 (Gene). The distinguishable units of information in a genotype that encode the phenotypical properties are called genes.

Definition 1.24 (Allele). An allele is a value of specific gene.

Definition 1.25 (Locus). The locus²⁸ is the position where a specific gene can be found in a genotype.

Figure 1.15 refines the relations of genotypes and phenotypes from the initial example for the spaces in Figure 1.13 by also marking genes, alleles, and loci. In the car customizing problem also mentioned earlier, the first gene could identify the color of the automobile. Its locus would then be 0 and it could have the alleles 00, 01, 10, and 11, encoding for red, white, green, and blue, for instance. The second gene (at locus 1) with the alleles 0 or 1 may define whether or not the car comes with climate control, and so on.

The Search Operations

In some problems, the search space $\mathbb{G}$ may be identical to the problem space $\mathbb{X}$. If we go back to our previous examples, we will find that there exist a lot of optimization strategies that work directly on vectors of real numbers. When minimizing a real function, we could use such an approach (Evolution Strategies, for instance, see Chapter 5 on page 191) and set $\mathbb{G} = \mathbb{X} = \mathbb{R}$. However, we will hardly find default search procedures that work on configurations of cars. Assume that such a configuration consists of k features, which can either be included or excluded from an offer to the customer. We can then search in the space of binary strings of this length $\mathbb{G} = \mathbb{B}^k = \{\mathbf{true}, \mathbf{false}\}^k$, which is exactly what genetic algorithms (discussed in Section 3.1 on page 113) do. By using their optimization capabilities, we do not need to mess with the search and selection techniques but can rely on well-researched standard operations.

Definition 1.26 (Search Operations). The set Op contains the search operations searchOp which can be used by an optimization algorithm in order to explore the search

²⁷ <http://en.wikipedia.org/wiki/Allele> [accessed 2007-07-03]

²⁸ http://en.wikipedia.org/wiki/Locus_%28genetics%29 [accessed 2007-07-03]

space $\mathbb{G}$. Search operations can be defined with different arities²⁹. Equation 1.23, for instance, denotes an n -ary operator, i. e., one with n arguments. The result of a search operation is one element of the search space.

$$\text{searchOp} : \mathbb{G}^n \mapsto \mathbb{G} \quad (1.23)$$

In the following chapters of this book, we will encounter many different new operators for search. An unary search operation would, for example, be the mutation in genetic algorithms (see Section 3.4.2). Optimization processes are often initialized by creating random genotypes – the results of a search operation with zero arity (no parameters).

Search operations often involve randomized numbers. In such cases, it makes no sense to reason about their results like $\exists g_1, g_2 \in \mathbb{G} : g_2 = \text{searchOp}(g_1) \wedge \dots$. Instead, we need to work with probabilities like $\exists g_1, g_2 \in \mathbb{G} : g_2 = P(\text{searchOp}(g_1)) > 0 \wedge \dots$. Based on Definition 1.26, we will use the notation $Op(x)$ for the application of any of the operations $\text{searchOp} \in Op$ to the genotype x . With $Op^k(x)$ we denote k successive applications of (possibly different) search operators. If the parameter x is left away, i. e., just Op^k is written, this chain has to start with a search operation with zero arity. In the style of Badea and Stanciu [41] and Skubch [42; 43], we now can define:

Definition 1.27 (Completeness). A set Op of search operations searchOp is *complete* if and only if every point g_1 in the search space $\mathbb{G}$ can be reached from every other point $g_2 \in \mathbb{G}$ by applying only operations $\text{searchOp} \in Op$.

$$\forall g_1, g_2 \in \mathbb{G} \Rightarrow \exists k \in \mathbb{N} : P(g_1 = Op^k(g_2)) > 0 \quad (1.24)$$

Definition 1.28 (Weak Completeness). A set Op of search operations searchOp is *weakly complete* if and only if every point g in the search space $\mathbb{G}$ can be reached by applying only operations $\text{searchOp} \in Op$. A weakly complete set of search operations hence includes at least one parameterless function.

$$\forall g \in \mathbb{G} \Rightarrow \exists k \in \mathbb{N} : P(g = Op^k) > 0 \quad (1.25)$$

If the set of search operations is not complete, there are points in the search space which cannot be reached. Then, we are probably not able to explore the problem space adequately and possibly will not find satisfyingly good solution.

Definition 1.29 (Adjacency (Search Space)). A point g_2 is adjacent to a point g_1 in the search space $\mathbb{G}$ if it can be reached by applying a single search operation $\text{searchOp} \in Op$ to g_1 . Notice that the adjacency relation is not necessarily symmetric.

$$\text{adjacent}(g_2, g_1) = \begin{cases} \text{true} & \text{if } \exists \text{searchOp} \in Op : P(\text{searchOp}(g_1) = g_2) > 0 \\ \text{false} & \text{otherwise} \end{cases} \quad (1.26)$$

The Connection between Search and Problem Space

If the search space differs from the problem space, a translation between them is furthermore required. In our car example we would need to transform the binary strings processed by the genetic algorithm to objects which represent the corresponding car configurations and can be processed by the objective functions.

Definition 1.30 (Genotype-Phenotype Mapping). The genotype-phenotype mapping (GPM, or ontogenic mapping [44]) $\text{gpm} : \mathbb{G} \mapsto \mathbb{X}$ is a left-total³⁰ binary relation (see Section 26.7 on page 413) which maps the elements of search space $\mathbb{G}$ to elements in the problem space $\mathbb{X}$.

$$\forall g \in \mathbb{G} \exists x \in \mathbb{X} : \text{gpm}(g) = x \quad (1.27)$$




Fig. 1.15: The relation of genome, genes, and the problem space.

The only hard criterion we impose on genotype-phenotype mappings in this book is left-totality, i. e., that they map each element of the search space to at least one solution candidate. They *may* be functional relations if they are deterministic. Although it is possible to create mappings which involve random numbers and, hence, cannot be considered to be functions in the mathematical sense of Section 26.7.1 on page 414. Then, Equation 1.27 would need to be rewritten to Equation 1.28.

$$\forall g \in \mathbb{G} \exists x \in \mathbb{X} : P(\text{gpm}(g) = x) > 0 \quad (1.28)$$

Genotype-phenotype mappings *should* further be surjective [45], i. e., relate at least one genotype to each element of the problem space. Otherwise, some solution candidates can never be found and evaluated by the optimization algorithm and there is no guarantee whether the solution of a given problem can be discovered or not. If a genotype-phenotype mapping is injective, which means that it assigns distinct phenotypes to distinct elements of the search space, we say that it is free from *redundancy*. There are different forms of redundancy, some are considered to be harmful for the optimization process, others have positive influence³¹. Most often, GPMs are not bijective (since they are neither necessarily injective nor surjective). Nevertheless, if a genotype-phenotype mapping is bijective, we can construct an inverse mapping $\text{gpm}^{-1} : \mathbb{X} \mapsto \mathbb{G}$.

$$\text{gpm}^{-1}(x) = g \Leftrightarrow \text{gpm}(g) = x \quad \forall x \in \mathbb{X}, g \in \mathbb{G} \quad (1.29)$$

Based on the genotype-phenotype mapping, we can also define an adjacency relation for the problem space, which, of course, is also not necessarily symmetric.

Definition 1.31 (Adjacency (Problem Space)). A point x_2 is adjacent to a point x_1 in the problem space $\mathbb{X}$ if it can be reached by applying a single search operation `searchOp` to their corresponding elements in the problem space.

$$\text{adjacent}(x_2, x_1) = \begin{cases} \text{true} & \text{if } \exists g_1, g_2 : x_1 = \text{gpm}(g_1) \wedge x_2 = \text{gpm}(g_2) \wedge \text{adjacent}(g_2, g_1) \\ \text{false} & \text{otherwise} \end{cases} \quad (1.30)$$

By the way, we now have the means to define the term *local optimum* clearer. The original Definition 1.8 only applies to single objective functions, but with the use of the adjacency relation `adjacent`, the prevalence criterion $\succ$, and the connection between the search space and the problem space `gpm`, we clarify it for multiple objectives.

²⁹ <http://en.wikipedia.org/wiki/Arity> [accessed 2008-02-15]

³⁰ See Equation 26.51 on page 413 to Section 26.7 on page 413 for an outline of the properties of binary relations.

³¹ See Section 1.4.4 on page 46 for more information.

Definition 1.32 (Local Optimum). A (local) optimum $x_l^* \in \mathbb{X}$ of a set of objective functions F function is not worse than all points adjacent to it.

$$\forall x_l^* \in \mathbb{G} \Rightarrow (\forall x \in \mathbb{X} : \text{adjacent}(x, x_l^*) \Rightarrow \overline{x \succ x_l^*}) \quad (1.31)$$

The Objective Space and Optimization Problems

After the appropriate problem space has been defined, the search space has been selected and a translation between them (if needed) was created, we are almost ready to feed the problem to a global optimization algorithm. The main purpose of such an algorithm obviously is to find as many elements as possible from the solution space – We are interested in the solution candidates with the best possible evaluation results. This evaluation is performed by the set F of n objective functions $f \in F$, each contributing one numerical value describing the characteristics of a solution candidate x .³²

Definition 1.33 (Objective Space). The objective space $\mathbb{Y}$ is the space spanned by the codomains of the objective functions.

$$F = \{f_i : \mathbb{X} \mapsto Y_i : 0 < i \leq n, Y_i \subseteq \mathbb{R}\} \Rightarrow \mathbb{Y} = Y_1 \times Y_2 \times \dots \times Y_n \quad (1.32)$$

The set F maps the elements x of the problem space $\mathbb{X}$ to the objective space $\mathbb{Y}$ and, by doing so, gives the optimizer information about their qualities as solutions for a given problem.

Definition 1.34 (Optimization Problem). An optimization problem is defined by a five-tuple $(\mathbb{X}, F, \mathbb{G}, Op, \text{gpm})$ specifying the problem space $\mathbb{X}$, the objective functions F , the search space $\mathbb{G}$, the set of search operations Op , and the genotype-phenotype mapping gpm . In theory, such an optimization problem *can* always be solved if Op is complete and the gpm is surjective.

Generic search and optimization algorithms find optimal elements if provided with an optimization problem defined in this way. Evolutionary algorithms, which we will discuss later in this book, are generic in this sense. Other optimization methods, like genetic algorithms for example, may be more specialized and work with predefined search spaces and search operations.

Fitness as a Relative Measure of Utility

When performing a multi-objective optimization, i. e., $n = |F| > 1$, the elements of $\mathbb{Y}$ are vectors in $\mathbb{R}^n$. In Section 1.2.2 on page 9, we have seen that such vectors cannot be compared directly and that we need some (comparative) measure for what is “good”. In many optimization techniques, especially in evolutionary algorithms, this measure is used to map the objective space to a subset $\mathbb{V}$ of the positive real numbers $\mathbb{R}^+$. For each solution candidate, this single real number represents its *fitness* as solution for the given optimization problem. The process of computing such a fitness value is often not solely depending on the absolute objective values of the solution candidates but also on those of the other phenotypes known. It could, for instance, be position of a solution candidate in the list of investigated elements sorted according to the Pareto relation. Hence, fitness values often only have a meaning inside the optimization process [35] and may change by time, even if the objective values stay constant. In deterministic optimization methods, the value of a heuristic function which approximates how many modifications we will have to apply to the element in order to reach a feasible solution can be considered as the fitness.

³² See also Equation 1.3 on page 9.

Definition 1.35 (Fitness). The fitness³³ value $v(x) \in \mathbb{V}$ of an element x of the problem space $\mathbb{X}$ corresponds to its utility as solution or its priority in the subsequent steps of the optimization process. The space spanned by all possible fitness values $\mathbb{V}$ is normally a subset of the positive real numbers $\mathbb{V} \subseteq \mathbb{R}^+$.

The origin of the term fitness has been borrowed biology³⁴ [46, 47] by the evolutionary algorithms community. When the first applications of genetic algorithms were developed, the focus was mainly on single-objective optimization. Back then, they called this single function fitness function and thus, set objective value $\equiv$ fitness value. This point of view is obsolete in principle, yet you will find many contemporary publications that use this notion. This is partly due the fact that in simple problems with only one objective function, the old approach of using the objective values directly as fitness, i. e., $v(x) = f(x) \forall x \in \mathbb{X}$, can sometimes actually be applied. In multi-objective optimization processes, this is not possible and fitness assignment processes like those which we are going to elaborate on in Section 2.3 on page 82 are applied instead.

In the context of this book, fitness is subject to minimization, i. e., elements with smaller fitness are better than those with higher fitness. Although this definition differs from the biological perception of fitness, it complies with the idea that optimization algorithms are to find the minima of mathematical functions (if nothing else has been stated).

Futher Definitions

In order to ease the discussions of different global optimization algorithms, we furthermore define the data structure *individual*. Especially evolutionary algorithms, but also many other techniques, work on sets of such individuals. Their fitness assignment processes determine fitness values for the individuals relative to all elements of these *populations*.

Definition 1.36 (Individual). An individual p is a tuple $(p.g, p.x)$ of an element $p.g$ in the search space $\mathbb{G}$ and its corresponding element $p.x$ in the problem space $\mathbb{X}$.

$$\forall \text{ individuals } p = (p.g, p.x) \in \mathbb{G} \times \mathbb{X} \Rightarrow p.x = \text{gpm}(p.g) \quad (\wedge p.g \in \mathbb{G} \wedge p.x \in \mathbb{X}) \quad (1.33)$$

Besides this basic individual structure, many practical realizations of optimization algorithms use such a data structure to store additional information like the objective and fitness values. Then, we will consider individuals as tuples in $\mathbb{G} \times \mathbb{X} \times Z$, where Z is the space of the additional information stored – $Z = \mathbb{Y} \times \mathbb{V}$, for instance. In the algorithm definitions later in this book, we will often access the phenotypes $p.x$ without explicitly using the genotype-phenotype mapping, since the relation of $p.x$ and $p.g$ complies to Equation 1.33 by definition.

Definition 1.37 (Population). A population Pop is a set of individuals used during an optimization process.

$$Pop \subseteq \mathbb{G} \times \mathbb{X} : \forall p = (p.g, p.x) \in Pop \Rightarrow p.x = \text{gpm}(p.g) \quad (1.34)$$

As already mentioned, the fitness $v(x)$ of an element x in the problem space $\mathbb{X}$ often not solely depends on the element itself. Normally, it is rather a relative measure putting the features of x into the context of a set of solution candidates x . We denote this by writing $v(x, X)$. It is also possible that the fitness involves the whole individual data, including the genotypic and phenotypic structures. We can denote this by writing $v(p, Pop)$.

³³ [http://en.wikipedia.org/wiki/Fitness_\(genetic_algorithm\)](http://en.wikipedia.org/wiki/Fitness_(genetic_algorithm)) [accessed 2008-08-10]

³⁴ http://en.wikipedia.org/wiki/Fitness_%28biology%29 [accessed 2008-02-22]

1.3.2 Fitness Landscapes and Global Optimization

A very powerful metaphor in global optimization is the fitness landscape³⁵. Like many other abstractions in optimization, fitness landscapes have been developed and extensively been researched by evolutionary biologists [48, 49, 50, 51]. Basically, they are visualizations of the relationship between the genotypes or phenotypes in a given population and their corresponding reproduction probability. The idea of such visualizations goes back to Wright [48], who used level contours diagrams in order to outline the effects of selection, mutation, and crossover on the capabilities of populations to escape local optimal configurations. Similar abstractions arise in many other areas [52], like in physics of disordered systems like spin-glasses [53, 54], for instance.

In Chapter 2, we will discuss evolutionary algorithms, which are optimization methods inspired by natural evolution. The evolutionary algorithm research community has widely adopted the fitness landscapes as relation between individuals and their objective values [55, 56]. Langdon and Poli [57]³⁶ explain that fitness landscapes can be imagined as a view on a countryside from far above. The height of each point is then analogous to its objective value. An optimizer can then be considered as a short-sighted hiker who tries to find the lowest valley or the highest hilltop. Starting from a random point on the map, she wants to reach this goal by walking the minimum distance.

As already mentioned, evolutionary algorithms were first developed as single-objective optimization methods. Then, the objective values were directly used as fitness and the “reproduction probability”, i. e., the chance of a solution candidate for being subject of further investigation, was proportional to them. In multi-objective optimization applications with more sophisticated fitness assignment and selection processes, this simple approach does not reflect the biological metaphor correctly anymore.

In the context of this book we will book, we therefore deviate from this view. Since it would possibly be confusing for the reader if we used a different definition for fitness landscapes than the rest of the world, we introduce the new term *problem landscape* and keep using the term *fitness landscape* in the traditional manner.

Definition 1.38 (Problem Landscape).

The problem landscape $\Phi : \mathbb{X} \times \mathbb{N} \mapsto [0, 1] \subset \mathbb{R}^+$ maps all the points x in a problem space $\mathbb{X}$ to the cumulative probability of reaching them until (inclusively) the τ^{th} evaluation of a solution candidate. The problem landscape thus depends on the optimization problem and on the algorithm applied in order to solve the problem.

$$\Phi(x, \tau) = P(x \text{ has been visited until the } \tau^{th} \text{ individual evaluation}) \quad \forall x \in \mathbb{X}, \tau \in \mathbb{N} \quad (1.35)$$

This definition of problem landscape is very similar to the performance measure definition used by Wolpert and Macready [58; 59] in their No Free Lunch Theorem which will be discussed in Section 1.4.9 on page 54. In our understanding, problem landscapes are not only closer to the original meaning of fitness landscapes in biology, they also have another advantage. According to this definition, all entities involved in an optimization process directly influence the problem landscape. The choice of the search operations in the search space $\mathbb{G}$, the way the initial elements are picked, the genotype-phenotype mapping, the objective functions, the fitness assignment process, and the way individuals are selected for further exploration all have impact on Φ . We can furthermore make the following assumptions about $\Phi x \tau$, since it is basically a some form of cumulative distribution function (see Definition 27.18 on page 421).

$$\Phi(x, \tau_1) \geq \Phi(x, \tau_2) \quad \forall \tau_1 < \tau_2 \wedge x \in \mathbb{X}, \tau_1, \tau_2 \in \mathbb{N} \quad (1.36)$$

$$0 \leq \Phi(x, \tau) \leq 1 \quad \forall x \in \mathbb{X}, \tau \in \mathbb{N} \quad (1.37)$$

³⁵ http://en.wikipedia.org/wiki/Fitness_landscape [accessed 2007-07-03]

³⁶ This part of [57] is also online available at http://www.cs.ucl.ac.uk/staff/W.Langdon/FOGP/intro_pic/landscape.html [accessed 2008-02-15].

Referring back to Definition 1.34, we can now also define what optimization algorithms are.

Definition 1.39 (Optimization Algorithm). An optimization algorithm is a transformation $(\mathbb{X}, F, \mathbb{G}, Op, \text{gpm}) \mapsto \Phi$ of an optimization problem $(\mathbb{X}, F, \mathbb{G}, Op, \text{gpm})$ to a problem landscape Φ that will find at least one *local* optimum x_l^* for each optimization problem $(\mathbb{X}, F, \mathbb{G}, Op, \text{gpm})$ with a weakly complete set of search operations Op and a surjective genotype-phenotype mapping gpm if granted infinite processing time and if such an optimum exists (see Equation 1.38).

$$\exists x_l^* \in \mathbb{X} : \lim_{\tau \rightarrow \infty} \Phi(x_l^*, \tau) = 1 \quad (1.38)$$

An optimization algorithm is characterized by

- the way it assigns fitness to the individuals,
- the ways it selects them for further investigation,
- the way it applies the search operations, and
- the way it builds and treats its state information.

The first condition in Definition 1.40, the completeness of Op , is mandatory because the search space $\mathbb{G}$ cannot be explored fully otherwise. If the genotype-phenotype mapping gpm is not surjective, there exist points in the problem space $\mathbb{X}$ which can never be evaluated. Only if both conditions hold, it is guaranteed that an optimization algorithm can find at least one local optimum.

The best optimization algorithm for a given problem $(\mathbb{X}, F, \mathbb{G}, Op, \text{gpm})$ is the one with the highest values of $\Phi(\mathbf{x}^*, \tau)$ for the optimal elements $\mathbf{x}^*$ in the problem space and for the lowest values of τ . It may be interesting that this train of thought indicates that finding the best optimization algorithm for a given optimization problem is, itself, a multi-objective optimization problem.

Definition 1.40 (Global Optimization Algorithm). Global optimization algorithms are optimization algorithms that employs measures that prevent convergence to local optima and increase the probability of finding a global optimum.

For a perfect global optimization algorithm (given an optimization problem with weakly complete search operations and a surjective genotype-phenotype mapping), Equation 1.39 would hold. In reality, it can be considered questionable whether such an algorithm can actually be built.

$$\forall x_1, x_2 \in \mathbb{X} : x_1 \succ x_2 \Rightarrow \lim_{\tau \rightarrow \infty} \Phi(x_1, \tau) > \lim_{\tau \rightarrow \infty} \Phi(x_2, \tau) \quad (1.39)$$

Let us now give a simple example for problem landscapes and how they are influenced by the optimization algorithm applied to them. Figure 1.16 illustrates one objective function, defined over a finite subset $\mathbb{X}$ of the two-dimensional real plane, which we are going to optimize. We use the problem space $\mathbb{X}$ also as search space $\mathbb{G}$, so we can do not need a genotype-phenotype mapping. For optimization, we will use a very simple hill climbing algorithm³⁷, which initially randomly creates one solution candidate uniformly distributed in $\mathbb{X}$. In each iteration, it creates a new solution candidate from the known one using an unary search operation. The old and the new candidate are compared, and the better one is kept. Hence, we do not need to differentiate between fitness and objective values. In the example, *better* means *has lower fitness*. In Figure 1.16, we can spot one local optimum x_l^* and one global optimum $\mathbf{x}^*$. Between them, there is a hill, an area of very bad fitness. The rest of the problem space exhibits a small gradient into the direction its center. The optimization algorithm will likely follow this gradient and sooner or later discover x_l^* or $\mathbf{x}^*$. The chances of x_l^* are higher, since it is closer to the center of $\mathbb{X}$.

³⁷ Hill climbing algorithms are discussed thoroughly in Chapter 8.




Fig. 1.16: An example optimization problem.

With this setting, we have recorded the traces of two experiments with 1.3 million runs of the optimizer (8000 iterations each). From these records, we can approximate the problem landscapes very good.

In the first experiment, depicted in Figure 1.17, we used a search operation $\text{searchOp}_1 : \mathbb{X} \mapsto \mathbb{X}$ which created a new solution candidate normally distributed around the old one. In all experiments, we had divided $\mathbb{X}$ in a regular lattice. $\text{searchOp}_2 : \mathbb{X} \mapsto \mathbb{X}$, used in the second experiment, the new solution candidates are direct neighbors of the old ones in this lattice. The problem landscape Φ produced by this operator is shown in Figure 1.18. Both operators are complete, since each point in the search space can be reached from each other point by applying them.

$$\text{searchOp}_1(x) \equiv (x_1 + \text{random}_n(), x_2 + \text{random}_n()) \quad (1.40)$$

$$\text{searchOp}_1(x) \equiv (x_1 + \text{random}_u(-1, 1), x_2 + \text{random}_u(-1, 1)) \quad (1.41)$$

In both experiments, the first probabilities of the elements of the search space of being discovered are very low, near to zero in the first few iterations. To put it precise, since our problem space is a 36×36 lattice, this probability is $1/36^2$ in the first iteration. Starting with the tenth or so iteration, small peaks begin to form around the places where the optima are located. These peaks grow

Well, as already mentioned, this idea of problem landscapes and optimization reflects solely the author's views. Notice also that it is not always possible to define problem landscapes for problem spaces which are *uncountable infinitely* large. Since the local optimum x_l^* at the center of the large basin and the gradient points straighter into its direction, it has a higher probability of being found than the global optimum $\mathbf{x}^*$. The difference between the two search operators tested becomes obvious starting with approximately the 2000th iteration. In the hill climber with the operator utilizing the normal distribution, the Φ value of the global optimum begins to rise farther and farther, finally surpassing the one of the local optimum. Even if the optimizer gets trapped in the local optimum, it will still eventually discover the global optimum and if we had run this experiment longer, the according probability would have converge to 1. The reason for this is that with the normal distribution, all points in the search space have a non-zero probability of being found from all other points in the search space. In other words, all elements of the search space are adjacent.

The operator based on the uniform distribution is only able to create points in the direct neighborhood of the known points. Hence, if an optimizer gets trapped in the local optimum, it can never escape. If it arrives at the global optimum, it will never discover the local one. In Fig. 1.18.12, we can see that $\Phi(x_l^*, 8000) \approx 0.7$ and $\Phi(\mathbf{x}^*, 8000) \approx 0.3$. One of the two points will be the result of the optimization process.




Fig. 1.17.1: $\Phi(x, 1)$




Fig. 1.17.2: $\Phi(x, 2)$




Fig. 1.17.3: $\Phi(x, 5)$




Fig. 1.17.4: $\Phi(x, 10)$




Fig. 1.17.5: $\Phi(x, 50)$




Fig. 1.17.6: $\Phi(x, 100)$




Fig. 1.17.7: $\Phi(x, 500)$




Fig. 1.17.8: $\Phi(x, 1000)$




Fig. 1.17.9: $\Phi(x, 2000)$




Fig. 1.17.10: $\Phi(x, 4000)$




Fig. 1.17.11: $\Phi(x, 6000)$
















Fig. 1.17.12: $\Phi(x, 8000)$

Fig. 1.17: The problem landscape of the example problem derived with searchOp_1 .

Fig. 1.18.1: $\Phi(x, 1)$ Fig. 1.18.2: $\Phi(x, 2)$ Fig. 1.18.3: $\Phi(x, 5)$ Fig. 1.18.4: $\Phi(x, 10)$ Fig. 1.18.5: $\Phi(x, 50)$ Fig. 1.18.6: $\Phi(x, 100)$ Fig. 1.18.7: $\Phi(x, 500)$ Fig. 1.18.8: $\Phi(x, 1000)$ Fig. 1.18.9: $\Phi(x, 2000)$ Fig. 1.18.10: $\Phi(x, 4000)$ Fig. 1.18.11: $\Phi(x, 6000)$ Fig. 1.18.12: $\Phi(x, 8000)$ Fig. 1.18: The problem landscape of the example problem derived with searchOp_2 .

From the example we can draw four conclusions:

- Optimization algorithms discover good elements with higher probability than elements with bad characteristics. Well, this is what they should do.
- The success of optimization depends very much on the way the search is conducted.
- It also depends on the time (or the number of iterations) the optimizer allowed to use.
- Hill climbing algorithms are no global optimization algorithms since they have no means of preventing getting stuck at local optima.

1.3.3 Gradient Descend

Definition 1.41 (Gradient). A gradient³⁸ of a scalar field $f : \mathbb{R}^n \mapsto \mathbb{R}$ is a vector field which points into the direction of the greatest increase of the scalar field. It is denoted by ∇f or $\text{grad}(f)$.

Optimization algorithms depend on some form of gradient in objective or fitness space in order to find good individuals. In most cases, the problem space $\mathbb{X}$ is not a vector space over the real numbers $\mathbb{R}$, so we cannot directly differentiate the objective functions with Nabla operator³⁹ ∇F . Generally, samples of the search space are used to approximate the gradient. If we compare to elements x_1 and x_2 of problem space and find $x_1 \succ x_2$, we can assume that there is some sort of gradient facing downwards from x_2 to x_1 . When descending this gradient, we can hope to find an x_3 with $x_3 \succ x_1$ and finally the global minimum.

1.3.4 Other General Features

There are some further common semantics and operations that are shared by most optimization algorithms. Many of them, for instance, start out by randomly creating some initial individuals which are then refined iteratively. Optimization processes which are not allowed to run infinitely have to find out when to terminate. In this section we define and discuss general abstractions for such commonalities.

Iterations

Global optimization algorithms often iteratively evaluate solution candidates in order to approach the optima. We distinguish between evaluations τ and iterations t .

Definition 1.42 (Evaluation). The value $\tau \in \mathbb{N}_0$ denotes the number of solution candidates for which the set of objective functions F has been evaluated.

Definition 1.43 (Iteration). An iteration⁴⁰ refers to one round in a loop of an algorithm. It is one repetition of a specific sequence of instruction inside an algorithm.

Algorithms are referred to as *iterative* if most of their work is done by cyclic repetition of one main loop. In the context of this book, an iterative optimization algorithm starts with the first step $t = 0$. The value $t \in \mathbb{N}_0$ is the index of the iteration currently performed by the algorithm and $t + 1$ refers to the following step. One example for iterative algorithm is Algorithm 1.1. In some optimization algorithms like genetic algorithms, for instance, iterations are referred to as *generations*.

There often exists a well-defined relation between the number of performed solution candidate evaluations τ and the index of the current iteration t in an optimization process: Many global optimization algorithms generate and evaluate a certain number of individuals per generation.

³⁸ <http://en.wikipedia.org/wiki/Gradient> [accessed 2007-11-06]

³⁹ <http://en.wikipedia.org/wiki/Del> [accessed 2008-02-15]

⁴⁰ <http://en.wikipedia.org/wiki/Iteration> [accessed 2007-07-03]

Termination Criterion

The termination criterion `terminationCriterion()` is a function with access to all the information accumulated by an optimization process, including the number of performed steps t , the objective values of the best individuals, and the time elapsed since the start of the process. With `terminationCriterion()`, the optimizers determine when they have to halt.

Definition 1.44 (Termination Criterion). When the termination criterion function `terminationCriterion() ∈ {true, false}` evaluates to **true**, the optimization process will stop and return its results.

Some possible criteria that can be used to decide whether an optimizer should terminate or not are [60, 61, 62, 63]:

- The user may grant the optimization algorithm a maximum computation time. If this time has been exceeded, the optimizer should stop. Here we should note that the time needed for single individuals may vary, and so will the times needed for iterations. Hence, this time threshold can sometimes not be abided exactly.
- Instead of specifying a time limit, a total number of iterations $\hat{t}$ or individual evaluations $\hat{\tau}$ may be specified. Such criteria are most interesting for the researcher, since she often wants to know whether a qualitatively interesting solution can be found for a given problem using at most a predefined number of samples from the problem space.
- An optimization process may be stopped when no improvement in the solution quality could be detected for a specified number of iterations. Then, the process most probably has converged to a (hopefully good) solution and will most likely not be able to make further progress.
- If we optimize something like a decision maker or classifier based on a sample data set, we will normally divide this data into a training and a test set. The training set is used to guide the optimization process whereas the test set is used to verify its results. We can compare the performance of our solution when fed with the training set to its properties if fed with the test set. This comparison helps us detect when most probably no further generalization can be achieved by the optimizer and we should terminate the process.
- Obviously, we can terminate an optimization process if it has already yielded a sufficiently good solution.

In practical applications, we can apply any combination of the criteria above in order to determine when to halt. How the termination criterion is tested in an iterative algorithm is illustrated in Algorithm 1.1.

Algorithm 1.1: Example Iterative Algorithm

Input: [implicit] `terminationCriterion()`: the termination criterion

Data: t : the iteration counter

```

1 begin
2    $t \leftarrow 0$ 
   // initialize the data of the algorithm
3   while terminationCriterion() do
4     // perform one iteration - here happens the magic
      $t \leftarrow t + 1$ 
5 end
```

Minimization

Many optimization algorithms have been developed for single-objective optimization in their original form. Such algorithms may be used for both, minimization or maximization. Without loss of generality we will present them as minimization processes since this is the most commonly used notation. An algorithm that maximizes the function f may be transformed to a minimization using $-f$ instead.

Note that using the prevalence comparisons as introduced in Section 1.2.7 on page 19, multi-objective optimization processes can be transformed into single-objective minimization processes. Therefore $x_1 \succ x_2 \Leftrightarrow \text{cmp}_F(x_1, x_2) < 0$.

Modeling and Simulating

While there are a lot of problems where the objective functions are mathematical expressions that can directly be computed, there exist problem classes far away from such simple function optimization that require complex models and simulations.

Definition 1.45 (Model). A model⁴¹ is an abstraction or approximation of a system that allows us to reason and to deduce properties of the system.

Models are often simplifications or idealization of real-world issues. They are defined by leaving away facts that probably have only minor impact on the conclusions drawn from them. In the area of global optimization, we often need two types of abstractions:

1. The models of the potential solutions shape the problem space $\mathbb{X}$. Examples are
 - programs in Genetic Programming, for example for the Artificial Ant problem,
 - construction plans of a skyscraper,
 - distributed algorithms represented as programs for Genetic Programming,
 - construction plans of a turbine,
 - circuit diagrams for logical circuits, and so on.
2. Models of the environment in which we can test and explore the properties of the potential solutions, like
 - a map on which the Artificial Ant will move which is driven by the evolved program,
 - an abstraction from the environment in which the skyscraper will be built, with wind blowing from several directions,
 - a model of the network in which the evolved distributed algorithms can run,
 - a physical model of air which blows through the turbine,
 - the model of an energy source the other pins which will be attached to the circuit together with the possible voltages on these pins.

Models themselves are rather static structures of descriptions and formulas. Deriving concrete results (objective values) from them is often complicated. It often makes more sense to bring the construction plan of a skyscraper to life in a simulation. Then we can test the influence of various wind strengths and directions on building structure and approximate the properties which define the objective values.

Definition 1.46 (Simulation). A simulation⁴² is the computational realization of a model. Whereas a model describes abstract connections between the properties of a system, a simulation realizes these connections.

Simulations are executable, live representations of models that can be as meaningful as real experiments. They allow us to reason if a model makes sense or not and how certain objects behave in the context of a model.

⁴¹ http://en.wikipedia.org/wiki/Model_%28abstract%29 [accessed 2007-07-03]

⁴² <http://en.wikipedia.org/wiki/Simulation> [accessed 2007-07-03]

1.4 Problems in Optimization

*“Like they want problems
You want ’em – you got ’em”*

— Method Man [64]

After reading the classification of optimization algorithms in Section 1.1.1, one might ask naïvely why there are so many different approaches. One possible answer is: Simply because there are so many different kinds of optimization tasks. Each of them puts different obstacles into the way of the optimizers and comes with own, characteristic difficulties. In this section we want to introduce the major complications that we may encounter during optimization tasks.

Some of subjects discussed in the following text concern global optimization in general (multi-modality and overfitting, for instance), others apply especially to nature-inspired approaches like genetic algorithms (epistasis and neutrality, for example). Neglecting even a single one them during the design or process of optimization can render the whole efforts invested useless, even if the most efficient optimization techniques are applied. By giving clear definitions and comprehensive introductions to these topics, we want to raise the awareness of scientists and users in the industry and hope to enable them to use optimization algorithms more efficiently.

1.4.1 Premature Convergence and Multimodality

Definition 1.47 (Convergence). An optimization algorithm has *converged* if it keeps on producing solutions from a “small”⁴³ subset of the problem space, possibly neglecting better solutions (in terms of fitness) in other areas.

Normally, global optimization algorithms will converge at some point in time. In nature, a similar phenomenon can be observed [65]. The *niche preemption principle* states that a niche in a natural environment tends to become dominated by a single species [66]. One of the problems in global optimization (and basically, also in nature) is that we generally cannot determine whether the best solution currently known is a local or a global optimum and convergence is acceptable. In other words, we are not able to say whether we can stop the optimization process, whether we should concentrate on refining our current optimum, or whether we should examine other parts of the search space instead. In the examples to the problem landscape, we already have encountered that problem. While the optimizer in first test was able to find the global optimum (see Figure 1.17), the second one (illustrated in Figure 1.18) prematurely converged to the local minimum. Premature convergence can, of course, only occur if there are multiple (local) optima, i. e., the problem is multimodal as depicted in Fig. 1.20.3.

Definition 1.48 (Multimodality). A set of objective functions F is multimodal, if it has multiple local or global optima. [67, 68].

The existence of multiple global optima is not problematic – if we find some but not all of them, this may be a pity but the optimization could still be considered as successful in many cases. More complicated is the occurrence of numerous local optima (see Definition 1.32).

Definition 1.49 (Premature Convergence). An optimization process has prematurely converged to a local optimum if it is no longer able to explore other parts of the search space than the currently examined area *and* there exists another region that contains a solution superior to the currently exploited one [69, 70].

⁴³ according to a suitable metric like numbers of mutations applied to a given solution

In Figure 1.19 we illustrate how a maximization process prematurely converges. An optimization algorithm will pass several local optima in objective space before reaching a good result. If it gets stuck on such an intermediate solution and cannot proceed to other points in search space anymore, we speak of premature convergence.




Fig. 1.19.1: Example 1: Maximization




Fig. 1.19.2: Example 2: Minimization

Fig. 1.19: Premature convergence in the objective space.

There are many features and parameter settings of optimization algorithms that influence the convergence behavior [71]. Some algorithms are capable of self-adaptation, allowing them to change their strategies or to modify their parameters depending on their current state. Such behavior often is implemented in order to speed up the optimization process, but may lead to premature convergence to local optima [72, 73, 74]. Restarting the optimization process at randomly chosen points in time is one crude yet sometimes effective measure against premature convergence which should not go unmentioned here.

Domino Convergence

The phenomenon of *domino convergence* has been brought to the light of attention by Rudnick [75] who researched it in the context of his BinInt problem [75, 76] which is discussed in Section 20.2.5. In principle, domino convergence occurs when the solution candidates have features which contribute to significantly different degrees to the total fitness. If these features are encoded in separate genes (or building blocks) in the genotypes, those are likely to be treated with different priorities, at least in randomized or heuristic optimization methods.

Building blocks with very strong positive influence on the objective values, for instance, will quickly be adopted by the optimization process (“converge”). During this time, the alleles of genes with a smaller contribution play no role. They do not come into play until the more “important” blocks have been accumulated. Rudnick [75] called this sequential convergence phenomenon *domino convergence* due to its resemblance with a row of falling domino stones [76] (see Section 20.2.5).

In the worst case, the contributions of the less salient genes may almost look like noise and they are not optimized at all. Such a situation is also an instance of premature convergence, since we cannot find the global optimum (which would involve optimal configurations of all blocks). Here, restarting the optimization process will not help because it will always turn out the same way. Example problems which often are likely to cause domino convergence

are the Royal Road and the aforementioned BinInt problem, which you can find discussed in Section 20.2.4 and Section 20.2.5, respectively.

Exploration vs. Exploitation

The operations that create new solutions from existing ones have also a very large impact on the speed of convergence [77, 78]. The step size in Evolution Strategy is a good example for this issue: setting it properly is very important and leads us to the problem of exploration versus exploitation which can be observed in the other areas of global optimization as well. There exists a vast body of research on the trade-off between exploration and exploitation that optimization algorithms have to face [79, 80, 81, 82, 83, 84].

Definition 1.50 (Exploration). Exploration in terms of optimization means finding new points in areas of the search space which have not been investigated yet.

Since computers have only with limited memory, this usually means most to drop already evaluated individuals. Exploration is the only means to find new and maybe better solution structures. Until such a new, good candidate is found – which is not guaranteed at all – the performance of the optimization process degenerates since individuals inferior to the already known ones are evaluated.

Definition 1.51 (Exploitation). Exploitation means trying to improve the currently known solution(s) by performing small changes which lead to new individuals very close to them.

This way, performance improvements can be achieved since the chance of finding better solutions similar to the already known individuals are often good. However, other, possibly better, solutions located in distant areas of the problem space will not be discovered by such minor refinements.

Almost all parts of optimization strategies can either be used for increasing exploitation or in favor for exploration. Unary search operations can be built that improve an existing solution in small steps, hence being exploitation operators (as is done in Memetic Algorithms, for instance). They can also be implemented in a way that introduces much randomness into the individuals, making them effectively exploration operations.

Selection operations⁴⁴ choose the set of individuals which will be investigated further. They can either return a small group of best individuals or a wide spread of existing solution candidates.

Optimization algorithms that favor exploitation have high convergence speed but run the risk of not finding the optimal solution and maybe get stuck at a local optimum. Algorithms that perform excessive exploration may find the global optimum but it will take them very long to do so. A good example for this dilemma is the Simulated Annealing algorithm discussed in Chapter 10 on page 217. It is often modified to a *simulated quenching* called form which favors exploitation but loses the guaranteed convergence to the optimum. Generally, optimization algorithms should employ at least one search operation of explorative character and at least one which is able to exploit good solution further.

Diversity

Exploitation and exploration are directly linked with *diversity*: exploration increases the diversity whereas exploitation works against it. Diversity preservation is a major concern in optimization because the loss of it can lead to premature convergence to a local optimum. Therefore, it has been studied by researchers from many domains and different points of view. Some of them are:

⁴⁴ Selection will be discussed in Section 2.4 on page 92.

- genetic algorithms: Oei et al. [85], and Ronald [86; 87],
- evolutionary algorithms: Bosman and Thierens [88; 89], Laumanns et al. [90], Mostaghim [91], Spieth et al. [92], and Singh and Deb [93]
- Genetic Programming: Gustafson [94], [95], Burke et al. [96; 97; 98], Brameier and Banzhaf [99], and de Jong et al. [100]
- Particle Swarm Optimization: Wilke et al. [101].

In evolutionary algorithms, special fitness assignment procedures like sharing or variety preservation (see Section 2.3.4 and Section 2.3.5) are used to keep the diversity of the population on a high level. Another, much simpler approach aiming for convergence prevention is introduced in Section 2.4.8.

In Figure 1.20, we have sketched a set of different types of fitness landscapes which we are going to discuss in this section. The small bubbles represent solution candidates. An arrow from one bubble to another means that the second individual is found by applying one search operation to the first solution candidate. The objective values are subject to minimization.

1.4.2 Ruggedness and Weak Causality

“You can’t handcuff what you can’t snuff.”

— Cage Kennylz [102]

Ruggedness

We have already stated that optimization algorithms generally depend on some form of gradient in the objective or fitness space. The objective functions should be continuous and exhibit low total variation⁴⁵, so the optimizer can descend the gradient easily. If the objective functions fluctuate, i. e., go up and down, it becomes more complicated for the optimizers to find the right directions to proceed into. The more rugged a function gets, the harder it becomes to optimize it. For short, one could say ruggedness is multi-modality plus steep ascends and descends in the fitness landscape. Examples for ruggedness are Kauffman’s fitness landscape (see Section 20.2.1), the p-Spin model discussed in Section 20.2.2, and Bergman and Feldman’s jagged fitness landscape [104].

Causality

During an optimization process, new points in the search space are created by the search operations. Generally we can assume that the genotypes are the input of the search operations correspond to phenotypes which have previously been selected. The chance of being selected for further investigation is usually the higher, the better or the more promising an individual is. Reversing this statement suggests that individuals which are passed to the search operations are likely to have a good fitness. Well, at least they should have a better fitness than the discarded solution candidates. Since the fitness of a solution candidate depends on its properties, we can assume that their properties were not so bad either. It should thus be possible for the optimizer to perform slight changes to these properties in order to find out whether they can be improved further. We already have learned this under the subject of exploitation. Normally, slight changes in the properties should also lead to small changes in the objective values and, hence, in the fitness of the solution candidate.

Definition 1.52 (Causality). The principle of strong causality (locality) proposed by Rechenberg [105; 106] states that small changes in an object lead to small changes in its behavior [105, 106, 107].

⁴⁵ http://en.wikipedia.org/wiki/Total_variation [accessed 2008-04-23]




Fig. 1.20.1: Best case




Fig. 1.20.2: Low Variation




Fig. 1.20.3: Multimodal




Fig. 1.20.4: Rugged




Fig. 1.20.5: Deceptive




Fig. 1.20.6: Neutral




Fig. 1.20.7: Needle-In-A-Haystack




Fig. 1.20.8: Nightmare

Fig. 1.20: Different possible properties of fitness landscapes.

This principle should not only hold for the search spaces and operations that we design for optimization, it also applies to natural genomes. The offspring resulting from sexual reproduction of two fish, for instance, has a different genome. It is far more probable that such changes manifest in a unique color pattern of the scales, for example, instead of leading to a totally different creature.

Apart from this straightforward, informal explanation here, causality has been investigated thoroughly in different fields of optimization, such as Evolution Strategy [105, 108], structure evolution [109, 110], Genetic Programming [111, 107, 112, 108], genotype-phenotype mappings [113], search operators [108], and evolutionary algorithms in general [114, 115, 108].

But where is the connection between rugged fitness landscapes and causality? In fitness landscapes with low (weak) causality, small changes in the solution candidates often lead to large changes in the objective values as outlined in Fig. 1.20.4. It then becomes harder to decide which region of the problem space to explore and the optimizer cannot find reliable gradient information to follow. A small modification of a very bad solution candidate may then lead to a new local optimum and the best solution candidate currently known may be surrounded by points that are inferior to all other tested individuals.

In general, we can say the lower the causality of an optimization problem, the more rugged its fitness landscape will be, and hence, the worse will optimizers perform [116]. This does not necessarily mean that we cannot find good solutions, but it may take very long to do so.

Fitness Landscape Measures

As measures for the ruggedness of a fitness landscape (or their general difficulty), many different metrics have been proposed. Wedge and Kell [117] and Altenberg [118] provide nice lists of them in their work⁴⁶, which we summarize here:

- Weinberger [119] introduced the autocorrelation function and the correlation length of random walks.
- The correlation of the search operators was used by Manderick et al. [120] in conjunction with the autocorrelation.
- Jones and Forrest [121], [122] proposed the fitness distance correlation (FDC), the correlation of the fitness of an individual and its distance to the global optimum. This measure has been extended by researchers such as Clergue et al. [123], [124].
- The probability that search operations create offspring fitter than their parents, as defined by Rechenberg [105] and Beyer [125] (and called evolvability by Altenberg [126]), will be discussed in Section 1.4.4 on page 44 in depth.
- Simulation dynamics have been researched by Altenberg [126] and Grefenstette [127].
- Another interesting metric is the fitness variance of formae (Radcliffe and Surry [128]) and schemas (Reeves and Wright [129]).
- The error threshold method from theoretical biology [130, 131] has been adopted Ochoa et al. [132] for evolutionary algorithms. It is the “critical mutation rate beyond which structures obtained by the evolutionary process are destroyed by mutation more frequently than selection can reproduce them” [132].
- The negative slope coefficient (NSC) by Vanneschi et al. [133; 134] may be considered as an extension of Altenberg’s evolvability measure.
- Davidor [135] uses the epistatic variance as a measure of utility of a certain representation in genetic algorithms. We discuss the issue of epistasis in Section 1.4.5.
- The genotype-fitness correlation (GFC) of Wedge and Kell [117] is a new measure for ruggedness in fitness landscape and has been shown to be a good guide for determining optimal population sizes in Genetic Programming.

⁴⁶ Especially the one of Wedge and Kell [117] is beautiful and far more detailed than this summary here.

Autocorrelation and Correlation Length

As example, let us take a look at the autocorrelation function as well as the correlation length of random walks [119]. Here we borrow its definition from Verel et al. [136]:

Definition 1.53 (Autocorrelation Function). Given a random walk $(x_i, x_{i+1}, \dots)$, the autocorrelation function ρ of an objective function f is the autocorrelation function of the time series $(f(x_i), f(x_{i+1}), \dots)$.

$$\rho(k, f) = \frac{E[f(x_i) f(x_{i+k})] - E[f(x_i)] E[f(x_{i+k})]}{D^2[f(x_i)]} \quad (1.42)$$

where $E[f(x_i)]$ and $D^2[f(x_i)]$ are the expected value and the variance of $f(x_i)$.

The correlation length $\tau = -\frac{1}{\log \rho(1, f)}$ measures how the autocorrelation function decreases and summarizes the ruggedness of the fitness landscape: the larger the correlation length, the lower the total variation of the landscape.

1.4.3 Deceptiveness

*“Become awake and chop the heads of these snakes
It’s better off than dying in a pit full of snakes
The snake inside yourself, that be the first head you take”*

— The RZA [137]

Another annoying feature of fitness landscapes is deceptiveness (or deceptivity). The gradient deceptive objective function leads the optimizer away from the optima, as illustrated in Fig. 1.20.5.

The term deceptiveness is mainly used in the genetic algorithm⁴⁷ community in the context of the Schema Theorem. Schemas describe certain areas (hyperplanes) in the search space. If an optimization algorithm has discovered an area with a better average fitness compared to other regions, it will logically focus on exploring this area. Thus, that the highly fit areas the optimization process converges to contain the true optimum. Objective functions where this is not the case are called deceptive [138, 139, 140].

If the information accumulated by an optimizer actually guides it away from the optimum, a search algorithm will perform worse than a random walk or an exhaustive enumeration method. This issue is known for a long time [141, 142, 143, 144] and has been subsumed under the No Free Lunch Theorem which we discuss in Section 1.4.9.

Example for deceptiveness are the ND fitness landscapes outlined in Section 20.2.3 and the Trap Functions (see Section 20.2.3).

1.4.4 Neutrality and Redundancy

Neutrality

Definition 1.54 (Neutrality). We consider the outcome of the application of a search operation to a solution candidate as neutral if it yields no change in the objective values [145].

⁴⁷ We are going to discuss genetic algorithms in Chapter 3 on page 113 and the Schema Theorem in Section 3.7 on page 123.

For all optimization algorithms, it is problematic when the best solution candidate currently found is situated on a plane of the fitness landscape, i. e., all adjacent solution candidates have the same objective values. As illustrated in Fig. 1.20.6, there is then no gradient information and thus, no direction into which the optimization algorithm can proceed in a systematic manner. From its point of view, each search operation will yield identically individuals. Furthermore, optimization algorithms usually maintain a list of the best individuals found, which will then eventually overflow or require pruning.

The degree of neutrality ν defines the fraction of neutral results among all possible products of the search operations applied to a specific genotype [145]. We can generalize this measure to areas G in the search space $\mathbb{G}$ by averaging over all its elements. Regions where ν is close to one are considered as *neutral*.

$$\forall g_1 \in \mathbb{G} \Rightarrow \nu(g_1) = \frac{|\{\forall g_2 : P(g_2 = Op(g_1)) > 0 \wedge F(\text{gpm}(g_2)) = F(\text{gpm}(g_1))\}|}{|\{\forall g_2 : P(g_2 = Op(g_1)) > 0\}|} \quad (1.43)$$

$$\forall G \subseteq \mathbb{G} \Rightarrow \nu(G) = \frac{1}{|G|} \sum_{\forall g \in G} \nu(g) \quad (1.44)$$

Evolvability

Another metaphor in global optimization borrowed from biological systems is evolvability⁴⁸ [146]. Wagner [147; 148] points out that this word has two uses in biology. According to Kirschner and Gerhart [149], a biological system is evolvable if it is able to generate heritable, selectable phenotypic variations. Such properties can then be evolved and changed by natural selection. In its second sense, a system is evolvable if it can acquire new characteristics via genetic change that help the organism(s) to survive and to reproduce. Theories about how the ability of generating adaptive variants has evolved have been proposed by Riedl [150], Altenberg [151], Wagner and Altenberg [152], and Bonner [153], amongst others. We can adopt the idea of evolvability for global optimization as follows.

Definition 1.55 (Evolvability). The evolvability of an optimization process in its current state defines how likely the search operations will lead to solution candidates with new (and eventually, better) objectives values.

The direct *probability of success* [105, 125], i. e., the chance that search operators produce offspring fitter than their parents, is also sometimes referred to as *evolvability* in the context of evolutionary algorithms [118, 126]. The link between evolvability and neutrality has been discussed by many researchers. The evolvability of neutral parts of a fitness landscapes depends on the optimization algorithm used. It is especially low for hill climbing and similar approaches, since the search operations cannot provide improvements (or even changes). The optimization process then degenerates to a random walk, as illustrated in Fig. 1.20.6. The work of Beaudoin et al. [154] on the ND fitness landscapes⁴⁹ furthermore shows that neutrality may “destroy” useful information such as correlation.

Recent research in molecular evolution, on the other hand, indicates that the majority of mutations have no selective influence [155, 156] and that the transformation from genotypes to phenotypes is a many-to-one mapping. Wagner [148] states that neutrality in natural genomes is beneficial if it concerns only a subset of the properties of the offspring of a solution candidate while allowing meaningful modifications of the others. Toussaint and Igel [157] even go as far as declaring it a necessity for self-adaptation. The theory of *punctuated equilibria*⁵⁰ in biology by Eldredge and Gould [158] states that species experience long periods of evolutionary inactivity which are interrupted by sudden, localized, and rapid

⁴⁸ <http://en.wikipedia.org/wiki/Evolvability> [accessed 2007-07-03]

⁴⁹ See Section 20.2.3 on page 285 for a detailed elaboration on the ND fitness landscape.

⁵⁰ http://en.wikipedia.org/wiki/Punctuated_equilibrium [accessed 2008-07-01]

phenotypic evolutions [159, 160]. It is assumed that the populations explore neutral layers⁵¹ during the time of stasis until, suddenly, a relevant change in a genotype leads to a better adapted phenotype [161] and reproduces quickly.

“Uh?”, you may think, “How does this fit together?” The key to differentiate “good” and “bad” neutrality is its degree ν in relation to the number of possible solutions maintained by the optimization algorithms. Smith et al. [162] have used illustrative examples similar to Figure 1.21 to show that a certain amount of neutral reproductions can foster the optimization progress. In Fig. 1.21.1, basically the same scenario of premature convergence as in Fig. 1.19.1 is depicted. The optimizer is drawn to a local optimum from which it cannot escape anymore. Fig. 1.21.2 shows that a little shot of neutrality could form a bridge to the global optimum. The optimizer now has a chance to escape the smaller peak if it is able to find and follow that bridge, i. e., the evolvability of the system has increased. If this bridge gets wider, as sketched in Fig. 1.21.3, this chance increases. Of course, if the bridge gets too wide, we’ll end up in a scenario like in Fig. 1.20.6, where the optimizer cannot find any direction. Furthermore, in this scenario we expect the neutral bridge to lead to somewhere useful, which is also not necessarily the case in reality.




Fig. 1.21: Possible positive influence of neutrality.

Recently, the idea of utilizing the processes of molecular⁵² and evolutionary⁵³ biology as complement to Darwinian evolution for optimization gains interest [163]. Scientists like Hu and Banzhaf [164] have begun to study the application of metrics such as the evolution rate of gene sequences [165, 166] to evolutionary algorithms. Here, the degree of neutrality (synonymous vs. non-synonymous changes) seems to play an important role.

Examples for neutrality in fitness landscapes are the ND family (see Section 20.2.3), the NKp and NKq models (discussed in Section 20.2.1), and the Royal Road (see Section 20.2.4). Another common instance of neutrality is bloat in Genetic Programming, which is outlined in Section 4.10.4 on page 188.

Neutral Networks

From the idea of neutral bridges between different parts of the search space as sketched by Smith et al. [162], we can derive the concept of neutral networks.

⁵¹ Or neutral networks, as introduced in Section 1.4.4.

⁵² http://en.wikipedia.org/wiki/Molecular_biology [accessed 2008-07-20]

⁵³ http://en.wikipedia.org/wiki/Evolutionary_biology [accessed 2008-07-20]

Definition 1.56 (Neutral Network). Neutral networks are equivalence classes K of elements of the search space $\mathbb{G}$ which map to elements of the problem space $\mathbb{X}$ with the same objective values and are connected by chains of applications of the search operators Op [145].

$$\begin{aligned} \forall g_1, g_2 \in \mathbb{G} : g_1 \in K(g_2) \subseteq \mathbb{G} \Leftrightarrow \exists k \in \mathbb{N}_0 : P(g_2 = Op^k(g_1)) > 0 \wedge \\ F(\text{gpm}(g_1)) = F(\text{gpm}(g_2)) \end{aligned} \quad (1.45)$$

Barnett [145] states that a neutral network has the *constant innovation property* if

1. the rate of discovery of innovations keeps constant for a reasonably large amount of applications of the search operations [167], and
2. if this rate is comparable with that of an unconstrained random walk.

Networks with this property may prove very helpful if they connect the optima in the fitness landscape. Stewart [168] utilizes neutral networks and the idea of punctuated equilibria in his *extrema selection*, a genetic algorithm variant that focuses on exploring individuals which are far away from the centroid of the set of currently investigated solution candidates (but have still good objective values). Then again, Barnett [169] showed that populations in genetic algorithm tend to dwell in neutral networks of high dimensions of neutrality regardless of their objective values, which (obviously) cannot be considered advantageous.

The convergence on neutral networks has furthermore been studied by Bornberg-Bauer and Chan [170], van Nimwegen et al. [171], [172], and Wilke [173]. Their results show that the topology of neutral networks strongly determines the distribution of genotypes on them. Generally, the genotypes are “drawn” to the solutions with the highest degree of neutrality ν on the neutral network Beaudoin et al. [154].

Redundancy

Definition 1.57 (Redundancy). Redundancy in the context of global optimization is a feature of the genotype-phenotype mapping and means that multiple genotypes map to the same phenotype, i. e., the genotype-phenotype mapping is not injective.

$$\exists g_1, g_2 : g_1 \neq g_2 \wedge \text{gpm}(g_1) = \text{gpm}(g_2) \quad (1.46)$$

The role of redundancy in the genome is as controversial as the one of neutrality. There exist many accounts for its positive influence on the optimization process. Shipman et al. [174], for instance, tried to mimic desirable evolutionary properties of RNA folding [156]. They developed redundant genotype-phenotype mappings using voting, Turing-machine like binary instructions, Cellular automata, and random Boolean network [49]. All four mappings induced neutral networks which proved beneficial for exploring the problem space. Especially the last approach provided particular good results [175]. Possibly converse effects like epistasis (see Section 1.4.5) arising from the new genotype-phenotype mappings have not been considered in this study.

Redundancy can have much impact on the explorability of the problem space. If we imagine a one-to-one mapping, the translation of a slightly modified genotype will always result in a different phenotype. If there exists a many-to-one mapping between genotypes and phenotypes, the search operations can create offspring genotypes different from the parent which still translate to the same phenotype. The optimizer may now walk along a path through this neutral network. If many genotypes along this path can be modified to different offspring, many new solution candidates can be reached [174].

Further experiments of Shipman et al. [176], [177] also indicate the positive effects of neutrality in the genotype-phenotype mapping. In the Cartesian Genetic Programming method, neutrality is explicitly introduced in order to increase the evolvability (see Section 4.7.4 on page 168) [178, 179]. Yet, Rothlauf [115] and Shackleton et al. [175] show that simple uniform redundancy is not necessarily beneficial for the optimization process and may even slow it down.

Summary

Different from ruggedness which is always bad for optimization algorithms, neutrality has many aspects that may further and such that hinder the process of finding good solutions. Generally we can state that degrees of neutrality ν approaching 1 degenerate optimization processes to random walks. Some forms of neutral networks, accompanied by low (nonzero) values of ν can improve the evolvability and, hence, increase the chance of finding good solutions.

Needle-In-A-Haystack

One of the worst cases of fitness landscapes is the *needle-in-a-haystack* problem (NIAH) sketched in Fig. 1.20.7, where the optimum occurs as *isolated* spike in a plane. In other words, small instances of extreme ruggedness combine with a general lack of information in the fitness landscape. Such optimization problems are extremely hard to solve and the optimization processes often will converge prematurely or take very long to find the global optimum. An example for such fitness landscapes is the all-or-nothing property often inherent to Genetic Programming of algorithms, as discussed in Section 4.10.3 on page 188.

1.4.5 Epistasis

In biology, *epistasis*⁵⁴ is defined as a form of interaction between different genes [180]. The term was originally coined by Bateson [181] and meant that one gene suppresses the phenotypical expression of another gene. In the context of statistical genetics, epistasis was originally called “epistacy” by Fisher [182]. According to Lush [183], the interaction between genes is epistatic if the effect on the fitness from altering one gene depends on the allelic state of other genes. This understanding of epistasis comes very close to another biological expression: *pleiotropy*⁵⁵ means that a single gene influences multiple phenotypic traits [184]. In global optimization, we generally do not make such fine-grained distinctions.

Definition 1.58 (Epistasis). Epistasis [185, 186, 187] in global optimization means that a change in one gene of a genotype, introduced by a search operation, for instance, leads to changes in multiple otherwise unrelated phenotypical properties.

We speak of minimal epistasis when every gene is independent of every other gene and of maximal epistasis when no proper subset of genes is independent of any other gene [188, 189]. Such tight relations violate the causality previously discussed, since changes in the phenotypes and the objective values resulting from changes in the genotypes should be small. In a genome with high epistasis, a modification in a genotype will alter multiple properties of the corresponding phenotype. Epistasis is mainly an aspect of the way we define the genome $\mathbb{G}$ and the genotype-phenotype mapping gpm . We should avoid it where possible [185].

As Figure 1.22 outlines, epistasis has a strong influence on many of the previously discussed problematic features. If one gene can “turn off” or affect the effects of many other genes, a modification of this gene will lead to a large change in the features of the phenotype. Hence, the causality will be weakened and ruggedness ensues in the fitness landscape. On the other hand, subsequent changes of the turned-off genes may have no influence on the phenotype at all, which would then increase the degree of neutrality in the search space.

There is an essential difference between epistasis and conflicting objectives in multi-objective optimization. Let us visit our example with the car configuration one more time. Objective functions can be the gasoline consumption the car, the power of its motor, and maybe the appeal of its design. It is quite natural that improvements in the second objective

⁵⁴ <http://en.wikipedia.org/wiki/Epistasis> [accessed 2008-05-31]

⁵⁵ <http://en.wikipedia.org/wiki/Pleiotropy> [accessed 2008-03-02]




Fig. 1.22: The influence of epistasis on the fitness landscape.

will lead to degeneration in the first one and vice versa, because the property “environment friendly” and “horse powers of the motor” are normally contradicting. Therefore, the relation of these two properties would not be regarded as epistatic in global optimization. If we change the gene encoding the color of the car and thereby influence the appeal of the car, we won’t expect any alteration in the other two objectives. Such changes can only occur if the color gene also influences the genes that specify the motor features to some degree. Such an influence on unrelated properties is an example for epistasis. A good genome or genotype-phenotype mapping should allow changing the color without interfering with other properties.

Naudts and Verschoren [190] have shown both, that deceptiveness does not occur in situations with low epistasis and also that objective functions with high epistasis are not necessarily deceptive on the example of length-two binary string genomes. A discussion about different shapes of fitness landscapes under the influence of epistasis is given by Beerenwinkel et al. [191].

1.4.6 Overfitting and Oversimplification

In all scenarios where optimizers evaluate some of the objective values of the solution candidates by using training data, two other phenomena with negative influence can be observed: overfitting and oversimplification. This is, for instance, the case when we use the data gathered in a customer survey to build classifiers to predict buying behavior. In this case, we would evaluate the utility of the solution candidates (classifiers) on the basis of training data extracted from the survey. Overfitting and oversimplification can also be observed when we use simulations to determine the objective values. Here, the different configurations and environments simulated represent the training data.

Overfitting

Definition 1.59 (Overfitting). Overfitting⁵⁶ is the emergence of an overly-complicated model (solution candidate) in an optimization process resulting from the effort to provide the best possible results for as much of the available training data as possible [192, 193, 194, 195].

A model (solution candidate) $m \in \mathbb{X}$ created with a finite set of training data is considered to be overfitted if a less complicated, alternative model $m' \in \mathbb{X}$, $m' \neq m$ exists which has a smaller error for the set of all possible (maybe even infinitely many) available or producible data samples. This model m' may, however, have a larger error in the training data.

⁵⁶ <http://en.wikipedia.org/wiki/Overfitting> [accessed 2007-07-03]

The phenomenon of overfitting is best known and can often be encountered in the field of artificial neural networks or in curve fitting⁵⁷ [196, 197, 198, 199, 200]. The latter means we that have a set A of n training data samples (x_i, y_i) and want to find a function f that represents these samples as good as possible, i. e., $f(x_i) = y_i \forall (x_i, y_i) \in A$. There exists exactly one polynomial⁵⁸ of the degree $n - 1$ that fits to each such training data and goes through all its points.⁵⁹ Hence, if we restrict ourselves to polynomial regression, we know that there exists exactly one global optimum, this single perfect fit. In spite of that, to be precise, there will also be an infinite number of polynomials with a higher degree than $n - 1$ that also match the sample data perfectly. Such results would be considered as overfitted. In Figure 1.23 we have sketched this problem. The function $f_1(x) = x$ shown in Fig. 1.23.2 has been sampled three times, as sketched in Fig. 1.23.1. There exists no polynomial of a degree of two or less that fits to these samples than f_1 . Optimizers, however, could also find overfitted polynomials of a higher degree such as f_2 that match to the data, as shown in Fig. 1.23.3. Here, f_2 plays the role of the overly complicated model m which will perform as good as the simpler model m when tested with the training sets, but will fail to deliver good results for all other input data.




Fig. 1.23.1: Three sample points of f_1 .

Fig. 1.23.2: $m' \equiv f_1(x) = x$.

Fig. 1.23.3: $m \equiv f_2(x)$.

Fig. 1.23: Overfitting due to complexity.

A very common cause for overfitting is noise in the sample data. There exists no measurement device for physical processes which delivers perfect results without error. Also surveys that represent the opinions of people to some certain topic or randomized simulations will exhibit variations from the true interdependencies of the observed entities. Hence, data samples A based on measurements will always contain some noise.

In Figure 1.24 we have sketched how such noise may lead to overfitted results. Fig. 1.24.1 illustrates a simple physical process obeying some quadratic equation. This process has been measured using some technical equipment and the 100 noisy samples depicted in Fig. 1.24.2 has been obtained. Fig. 1.24.3 shows a function resulting from an optimization that fits perfectly to the data. It could, for instance, be a polynomial of degree 99 that goes right through all the points and thus, has an error of zero. Again, this complicated model, although being a perfect match to the measurements, does not accurately represent the physical law that produced the sample data and will not deliver precise results for new, different inputs. We give a short overview on the types of noise that should be considered during optimization in Section 1.4.7.

From the examples we can see that the major problem that results from overfitted solutions is the loss of generality.

⁵⁷ We will discuss overfitting in conjunction with Genetic Programming-based symbolic regression in Section 22.1 on page 341.

⁵⁸ <http://en.wikipedia.org/wiki/Polynomial> [accessed 2007-07-03]

⁵⁹ http://en.wikipedia.org/wiki/Polynomial_interpolation [accessed 2008-03-01]




Fig. 1.24.1: The original physical process.




Fig. 1.24.2: The measurement/training data.




Fig. 1.24.3: The overfitted result.

Fig. 1.24: Fitting noise.

Definition 1.60 (Generality). A solution of an optimization process is general if it is not only valid for the sample inputs $x_1, x_2, \dots, x_n$ which were used for training during the optimization process, but also for different inputs $x \neq x_i \forall i : 0 < i \leq n$ if such inputs x exist. Hence, a solution is also general if it is valid for all possible inputs.

Overfitting Prevention

There exist multiple techniques that can help us to prevent overfitting to a certain degree. It is most efficient to always apply multiple such techniques together in order to achieve best results.

Restriction of the Problem Space We can restrict the problem space $\mathbb{X}$ in a way that only solutions up to a given maximum complexity can be found.

Additional Optimization Criteria Furthermore, the functional objective function which solely concentrates on the error of the solution candidates should be augmented by penalty terms and secondary objective functions putting pressure into the direction of small and simple models Dietterich [195], Kearns et al. [201].

Much Training Data Large sets of sample data, although slowing down the optimization process, may improve the generalization capabilities of the solutions derived.

Changing Training Data/Simulation Scenarios If arbitrary many training datasets or training scenarios (for simulations, for instance) can be generated, there exist two approaches that prevent overfitting:

1. Use a new set of (randomized) scenarios for each evaluation of a solution candidate. Here, the resulting objective values may differ largely even if the same individual is evaluated twice in a row, introducing some incoherence and ruggedness into the fitness landscape.
2. At the beginning of each iteration of the optimizer, generate a new same set of (randomized) scenarios which is used for all individual evaluations during that iteration.

In both cases, it is important to use more than one training sample or scenario per evaluation and set the resulting objective value to the average (or better median) of the outcomes. Otherwise, the fluctuations of the objective values between the iterations will be very large, making it hard for the optimizers to follow a stable gradient for multiple steps.

Early Stopping A very simple method to prevent overfitting is to limit the runtime of the optimizers [192]. It is commonly assumed that learning processes normally first find relatively general solutions which subsequently begin to overfit because also the noise “is learned”.

Decay of Influence For the same reason, some algorithms allow to decrease the rate at which the solution candidates are modified by time. Such a decay of the learning rate makes overfitting less likely.

Dividing Data into Training and Test Sets If only one finite set of data samples A is available for training/optimization, it is common practice to separate it into a set of training data A_t and a set of test cases A_c . During the optimization process, only the training data A_t is used. The resulting solution are tested with the test cases A_c afterwards. If their behavior is significantly worse when applied to A_c than when applied to A_t , they are probably overfitted.

We can use the same approach to detect when we should stop the optimization process. The best known solution candidates can be checked with the test cases each iteration, without influencing their objective values which solely depend on the training data. If their performance on the test cases begins to decrease, we should stop.

Oversimplification

Oversimplification (also called overgeneralization) is the opposite of overfitting. Whereas overfitting denotes the emergence of overly-complicated solution candidates, oversimplified solutions are not complicated enough. Although they represent the training samples used during the optimization process seemingly well, they are rough overgeneralizations which fail to provide good results for cases not part of the training.

A very common cause for oversimplification is sketched in Figure 1.25: Training sets that only represent a fraction of the set of possible inputs. While this is normally the case, one should always be aware that such an incomplete coverage may fail to represent some of the dependencies and characteristics of the data, which then may lead to oversimplified solutions.

Fig. 1.25.1 shows a cubic function. Since it is a polynomial of degree three, four sample points are needed for its unique identification. We, however, have sampled it only three times, as shown in Fig. 1.25.2. By doing so, we lose some of vital characteristics of the function. Fig. 1.25.3 depicts a square function – the polynomial of the lowest degree that we can fit exactly to these samples. Although it is a perfect match, this function does not fit to any other point on the original cubic curve.




Fig. 1.25.1: The “real system”.




Fig. 1.25.2: The measurement/training data.




Fig. 1.25.3: The oversimplified result.

Fig. 1.25: Oversimplification.

1.4.7 Robustness and Noise

Robustness

In global optimization, we always seek the global optima of the objective functions. Well, at least from the theoretical point of view. From the practical point of view, we not necessarily

do. One reason for this is that the solutions of practical problems often are parameters or plans to be realized as components or to be put into action by human beings or machines. This can only be done if they allow a certain degree of imprecision. There exists no process in the world that is 100% accurate. Thus, we would sometimes favor a local optimum (or even non-optimal elements) for which slight deviations also lead to gentle performance degenerations over a global optimum located in a highly rugged area of the search space. [202]. In other words, local optima in regions of the search space with strong causality are sometimes better than global optima with weak causality. Of course, the level of acceptability is application-dependent. Figure 1.26 illustrates the issue of local optima which are robust vs. global optima which are not. More examples from the real world are:

- If we want to optimize the control parameters of an airplane or a nuclear power plant, we certainly do not want to use the global optima if a slight perturbation can have hazardous effects on the system. [203]
- Wiesmann et al. [204; 205] bring up the topic of manufacturing tolerances in multilayer optical coatings. It is no use to find optimal configurations if they only perform optimal when manufactured to a precision which is either impossible or too hard to achieve on a constant basis.
- We could optimize the decision process on which roads should be precautionary salted for areas with marginal winter climate. The global optimum of this problem is likely to depend on the daily (or even current) weather forecast and may therefore be constantly changing. Handa et al. [206] point out that it is practically infeasible to let road workers follow a constantly changing plan and circumvent this dynamic robustness issue.
- Tsutsui et al. [203], [207] even find a nice analogue in nature: The phenotypic characteristics of an individual are described by its genetic code. Nevertheless, during the interpretation of this code, perturbations like abnormal temperature, nutritional imbalances, injuries, illnesses and so on may occur. If the phenotypic features emerging under this influence have low fitness, the organism cannot survive and procreate. Thus, even a species with good genotypic material will die out if its phenotypic features become too sensitive to perturbations. Species robust against them, on the other hand, will survive and evolve.




Fig. 1.26: A robust local optimum vs. a “unstable” global optimum.

For the special case where the problem space corresponds to the real vectors ($\mathbb{X} \subseteq \mathbb{R}^n$), several approaches for dealing with the problem of robustness have been developed. Inspired by Taguchi methods⁶⁰ [208], for example, Greiner [209; 210] developed an approach

⁶⁰ http://en.wikipedia.org/wiki/Taguchi_methods [accessed 2008-07-19]

where possible disturbances are represented by a vector $\boldsymbol{\delta} = (\delta_1, \delta_2, \dots, \delta_n)^T$, $\delta_i \in \mathbb{R}$. If the distribution and influence of the δ_i are known, the objective function $f(\mathbf{x}) : \mathbf{x} \in \mathbb{X}$ can be rewritten as $\tilde{f}(\mathbf{x}, \boldsymbol{\delta})$ [204]. If $\boldsymbol{\delta}$ is normally distributed, we could, for instance, write $\tilde{f}\left((x_1 + \delta_1, x_2 + \delta_2, \dots, x_n + \delta_n)^T\right)$. It would then make sense to sample the probability distribution of $\boldsymbol{\delta}$ a number of t times and to use the mean values of $\tilde{f}(\mathbf{x}, \boldsymbol{\delta})$ for each objective value evaluation during the optimization process. For cases where the optimal value y of the objective function f is known, we can minimize Equation 1.47 and, by doing so, turn our optimization algorithm into a maximum likelihood estimator as shown in Section 27.6.1 and according to Equation 27.252 on page 453 (this approach is also used in the work of Wiesmann et al. [205; 204]).

$$f'(\mathbf{x}) = \frac{1}{t} \sum_{i=1}^t \left(y - \tilde{f}(\mathbf{x}, \boldsymbol{\delta}_i) \right)^2 \quad (1.47)$$

Noise

In general, we can distinguish three types of noise in optimization: First, *(i)* there is noise in the training data that we use as basis for learning and that may cause *overfitting*. This noise results from the fact that no measurement is 100% accurate and always exists when we try to fit a model to measured data. The second form of noise subsumes the perturbations that are likely to occur in the subsequent process in which our solutions are used – it is the cause why we need the best *robust* solutions and not just the globally optimal ones in many cases. We can subdivide this category into perturbations that may arise from inaccuracies in *(ii)* the process of realizing our solutions and *(iii)* environmentally induced perturbations during the applications of the products.

We can visualize this when thinking about the process of developing the perfect tire for a car. Well, as input we would have all sorts of shiny coefficients, material, and geometric constants measured from all known types of wheels and rubber available. In these constants, there surely is already a certain degree of noise and imprecision included *(i)*. During the process of optimization, we will simulate the behavior of the solution candidates for finding out their utility. The result will be the construction plan for a perfect tire, likely incorporating different materials and structures. Surely we would hope that the tires created according to the plan will not fall apart if we, accidentally, use 0.0001% too much of a specific rubber component *(ii)*. Also, it should not behave unexpectedly when used in scenarios different from our simulations *(iii)* and instead be applicable in all driving scenarios likely to occur.

Many global optimization algorithms and theoretical results have been proposed which can deal with noise, some of them are:

- genetic algorithms [211, 203, 207, 212, 213, 214],
- Evolution Strategies [215, 216, 217],
- Particle Swarm Optimization [218, 219],
- **TODO**

The effects of noise in optimization have been studied by various researchers, Miller and Goldberg [220; 221], Lee and Wong [222], and Gurin and Rastrigin [223] are some of them.

1.4.8 Dynamically Changing Fitness Landscape

It should also be mentioned that there exist problems with dynamically changing fitness landscapes [224, 225, 226, 227]. The task of an optimization algorithm is then to provide solution candidates with momentarily optimal objective values for each point in time. Here we have the problem that an optimum in iteration t will possibly not be an optimum in generation $t + 1$ anymore.

Problems with dynamic characteristics can for example be tackled with special forms [228] of

- evolutionary algorithms [229, 230, 231, 232, 233, 234, 235],
- genetic algorithms [236, 237, 238, 239, 214],
- Particle Swarm Optimization [240, 241, 242, 243, 244],
- Differential Evolution [245, 246], and
- Ant Colony Optimization [247, 248]

The moving peaks benchmarks by Branke [227] and Morrison and De Jong [225] are good examples for dynamically changing fitness landscapes. You can find them discussed in Section 20.1.3 on page 280.

1.4.9 No Free Lunch Theorem

By now, we know the most important problems that we may encounter when we want to apply an optimization algorithm to a given problem. Furthermore, we have seen that we can even dispute what actually an optimum is if we optimize multiple criteria at once. The fact that there is most likely no optimization method that can outperform all others on all problems can, thus, easily be accepted. Instead, there exists a variety of optimization methods specialized in solving different types of problems. There are also algorithms which deliver good results for a many different problem classes, but may be outperformed by highly specialized methods in each of them.

These facts have been formalized by Wolpert and Macready [58; 59] in their *No Free Lunch Theorems*⁶¹ (NFL) for search and optimization algorithms. The first step in proofing that optimization algorithms perform differently is to find a measure for performance. Wolpert and Macready [59] focus on single-objective optimization therefore use a measure which is very similar to the problem landscapes introduced in Definition 1.38 on page 29. The value of the problem landscape $\Phi(x, \tau)$ is the cumulative probability that the optimizer has visited the element $x \in \mathbb{X}$ until (inclusively) the τ^{th} evaluation of the objective function(s). Wolpert and Macready [59] include the algorithm a and the (single) objective function f into this notation and write $\Phi(x, \tau, a, f)$ in order to emphasize that the value Φ depends on those two too. In particular, they prove that the sum of these problem landscape values over all possible objective functions f is always identical for all optimization algorithms. For two optimizers a_1 and a_2 , this means that

$$\sum_{\forall f} \Phi(x, \tau, a_1, f) = \sum_{\forall f} \Phi(x, \tau, a_2, f) \quad \forall \tau \in \mathbb{N}, x \in \mathbb{X} \quad (1.48)$$

From this theorem, we can immediately follow that, in order to outperform a_1 in one optimization problem, a_2 will necessarily perform worse in another. Figure 1.27 visualizes this issue. It shows that general optimization algorithms like evolutionary algorithms can solve a variety of problem class with reasonable performance. We need to remember that the all non-zero values of a performance measure Φ mean that a problem can be solved by the algorithm. The higher these values, the faster will the problem be solved. hill climbing approaches, for instance, will be much faster than evolutionary algorithms if the objective functions are steady and monotonous. Greedy search methods will perform fast on all problems with matroid⁶² structure. Evolutionary algorithms will most often still be able to solve these problems, it just takes them longer to do so. The performance of hill climbing and greedy approaches degenerates in other classes of optimization tasks as a trade-off for their high utility in their “area of expertise”.

⁶¹ http://en.wikipedia.org/wiki/No_free_lunch_in_search_and_optimization [accessed 2008-03-28]

⁶² <http://en.wikipedia.org/wiki/Matroid> [accessed 2008-03-28]




Fig. 1.27: A visualization of the No Free Lunch Theorem.

One interpretation of the No Free Lunch Theorem is that it is impossible for any optimization algorithm to outperform all random walks or exhaustive enumerations on all possible problems. For every problem where a given method leads to good results, we can construct a problem where the same method has exactly the opposite effect.

Another interpretation is that every working optimization algorithm utilizes some form of problem specific knowledge. Radcliffe [249] states that without such knowledge, search algorithms cannot exceed the performance of simple enumerations. Incorporating knowledge starts with relying on as simple assumptions as “if x is a good solution candidate, then we can expect other good solution candidates in its vicinity”, i. e., strong causality. The more (correct) problem specific knowledge is integrated (correctly) into the algorithm structure, the better will the algorithm perform. On the other hand, knowledge correct for one class of problems is possibly totally incorrect for another class. In reality, we use optimizers to solve a given sort of problems and are not interesting in their performance when (wrongly) applied to other classes.

Today, there exists a wide range of work on No Free Lunch Theorems for many different aspects of machine learning. The website <http://www.no-free-lunch.org/> [accessed 2008-03-28] gives a good overview about them. Further summaries and extensions have been provided by Köppen et al. [250] and Igel and Toussaint [251; 252]. Radcliffe and Surry [45] discusses the NFL in the context of evolutionary algorithms and the representations used as search spaces.

1.5 Formae and Search Space/Operator Design

Most global optimization algorithms share the premise that solutions to problems are either elements of a somewhat continuous space that can be approximated stepwise or that they can be composed of smaller modules which have good attributes even when occurring separately.

The design of the search space (or genome) $\mathbb{G}$ and the genotype-phenotype mapping gpm is vital for the success of the optimization process. It determines to what degree these expected features can be exploited by defining how the properties and the behavior of the solution candidates are encoded and how the search operations influence them. In this chapter, we will first discuss a general theory about how properties of individuals can be defined, classified, and how they are related. We will then outline some general rules for the design of the genome which are inspired by our previous discussion of the possible problematic aspects of fitness landscapes.

1.5.1 Forma Analysis

The Schema Theorem has been stated for genetic algorithms by Holland [253] in its seminal work [253, 254, 80]. In this section, we are going to discuss it in the more general version from Weicker [255] as introduced by Radcliffe and Surry [128] and Surry [256] in [257, 249, 258, 258, 128].

The different individuals p in the population Pop of the search and optimization algorithms are characterized by their properties ϕ . Whereas the optimizers themselves focus mainly on the phenotypical properties since these are evaluated by the objective functions, the properties of the genotypes may be of interest in an analysis of the optimization performance.

A rather structural property ϕ_1 of formulas $f : \mathbb{R} \mapsto \mathbb{R}$ in symbolic regression⁶³ would be whether it contains the mathematical expression $x+1$ or not. We can also declare a behavioral property ϕ_2 which is **true** if $|f(0) - 1| \leq 0.1$ holds, i. e., if the result of f is close to a value 1 for the input 0, and **false** otherwise. Assume that the formulas were decoded from a binary search space $\mathbb{G} = \mathbb{B}^n$ to the space of trees that represent mathematical expression by a genotype-phenotype mapping. A genotypical property then would be if a certain sequence of bits occurs in the genotype $p.g$ and a phenotypical property is the number of nodes in the phenotype $p.x$, for instance. If we try to solve a graph-coloring problem, for example, a property $\phi_3 \in \{\text{black}, \text{white}, \text{gray}\}$ could denote the color of a specific vertex q as illustrated in Figure 1.28.




Fig. 1.28: An graph coloring-based example for properties and formae.

⁶³ More information on symbolic regression can be found in Section 22.1 on page 341.

In general, we can imagine the properties ϕ_i to be some sort of functions that map the individuals to property values. ϕ_1 and ϕ_2 would then both map the space of mathematical functions to the set $\mathbb{B} = \{\text{true}, \text{false}\}$ whereas ϕ_3 maps the space of all possible colorings for the given graph to the set $\{\text{white}, \text{gray}, \text{black}\}$. On the basis of the properties ϕ_i we can define equivalence relations⁶⁴ $\sim_{\phi_i}$:

$$p_1 \sim_{\phi_i} p_2 \Rightarrow \phi_i(p_1) = \phi_i(p_2) \quad \forall p_1, p_2 \in \mathbb{G} \times \mathbb{X} \quad (1.49)$$

Obviously, for each two solution candidates and x_1 and x_2 , either $x_1 \sim_{\phi_i} x_2$ or $x_1 \not\sim_{\phi_i} x_2$ holds. These relations divide the search space into equivalence classes $A_{\phi_i=v}$.

Definition 1.61 (Forma). An equivalence class $A_{\phi_i=v}$ that contains all the individuals sharing the same characteristic v in terms of the property ϕ_i is called a *forma* [258] or *predicate* [259].

$$A_{\phi_i=v} = \{\forall p \in \mathbb{G} \times \mathbb{X} : \phi_i(p) = v\} \quad (1.50)$$

$$\forall p_1, p_2 \in A_{\phi_i=v} \Rightarrow p_1 \sim_{\phi_i} p_2 \quad (1.51)$$

The number of formae induced by a property, i. e., the number of its different characteristics, is called its *precision* [258]. The precision of ϕ_1 and ϕ_2 is 2, for ϕ_3 it is 3. We can define another property $\phi_4 \equiv f(0)$ denoting the value a mathematical function has for the input 0. This property would have an uncountable infinite large precision.




Fig. 1.29: Example for formae in symbolic regression.

Two formae $A_{\phi_i=v}$ and $A_{\phi_j=w}$ are said to be *compatible*, written as $A_{\phi_i=v} \bowtie A_{\phi_j=w}$, if there can exist at least one individual which is an instance of both.

$$A_{\phi_i=v} \bowtie A_{\phi_j=w} \Leftrightarrow A_{\phi_i=v} \cap A_{\phi_j=w} \neq \emptyset \quad (1.52)$$

$$A_{\phi_i=v} \bowtie A_{\phi_j=w} \Leftrightarrow \exists p \in \mathbb{G} \times \mathbb{X} : p \in A_{\phi_i=v} \wedge p \in A_{\phi_j=w} \quad (1.53)$$

$$A_{\phi_i=v} \bowtie A_{\phi_i=w} \Rightarrow w = v \quad (1.54)$$

⁶⁴ See the definition of equivalence classes in Section 26.7.3 on page 416.

Of course, two different formae of the same property ϕ_i , i. e., two different characteristics of ϕ_i , are always incompatible. In our initial symbolic regression example hence $A_{\phi_1=\text{true}} \not\bowtie A_{\phi_1=\text{false}}$ since it is not possible that a function f contains a term $x + 1$ and at the same time does not contain it. All formae of the properties ϕ_1 and ϕ_2 on the other hand are compatible: $A_{\phi_1=\text{false}} \bowtie A_{\phi_2=\text{false}}$, $A_{\phi_1=\text{false}} \bowtie A_{\phi_2=\text{true}}$, $A_{\phi_1=\text{true}} \bowtie A_{\phi_2=\text{false}}$, and $A_{\phi_1=\text{true}} \bowtie A_{\phi_2=\text{true}}$ all hold. If we take ϕ_4 into consideration, we will find that there exist some formae compatible with some of ϕ_2 and some that are not, like $A_{\phi_2=\text{true}} \bowtie A_{\phi_4=1}$ and $A_{\phi_2=\text{false}} \bowtie A_{\phi_4=2}$, but $A_{\phi_2=\text{true}} \not\bowtie A_{\phi_4=0}$ and $A_{\phi_2=\text{false}} \not\bowtie A_{\phi_4=0.95}$.

The discussion of forma and their dependencies stems from the evolutionary algorithm community and there especially from the supporters of the Building Block Hypothesis. The idea is that the algorithm first discovers formae which have a good influence on the overall fitness of the solution candidates. The hope is that there are many compatible ones under these formae that are then gradually combined in the search process.

In this text we have defined formae and the corresponding terms on the basis of individuals p which are records that assign an element of the problem spaces $p.x \in \mathbb{X}$ to an element of the search space $p.g \in \mathbb{G}$. Generally, we will relax this notation and also discuss forma directly in the context of the search space $\mathbb{G}$ or problem space $\mathbb{X}$, when appropriate.

1.5.2 Genome Design

In software engineering, there are some design patterns⁶⁵ that describe good practice and experience values. Utilizing these patterns will help the software engineer to create well-organized, extensible, and maintainable applications.

Whenever we want to solve a problem with global optimization algorithms, we need to define the structure of a genome. The individual representation along with the genotype-phenotype mapping is a vital part of genetic algorithms and has major impact on the chance of finding good solutions.

We have already discussed the basic problems that we may encounter during optimization. The choice of the search space, the search operations, and the genotype-phenotype mapping have major impact on the chance of finding good solutions. After formalizing the ideas of properties and formae, we will now outline some general best practices for the genome design from different perspectives. These principles can lead to finding better solutions or higher optimization speed if considered in the design phase [115, 260].

In Goldberg [139] defines two general design patterns for genotypes in genetic algorithm which we will state here in the context of the forma analysis [260]:

1. The representations of the formae in the search space should be as short as possible and the representations of different, compatible phenotypic formae should not influence each other.
2. The alphabet of the encoding and the lengths of the different genes should be as small as possible.

Both rules target for minimal redundancy in the genomes. We have already mentioned in Section 1.4.4 on page 46 that uniform redundancy slows down the optimization process. Especially the second rule focuses on this cause of neutrality by discouraging the use of unnecessary large alphabets for encoding in a genetic algorithm. Palmer and Kershenbaum [261; 262] define additional rules for tree-representations in [261, 263], which have been generalized by Nguyen [260]:

3. A good search space and genotype-phenotype mapping should be able to represent all phenotypes, i. e., be surjective (see Section 26.7 on page 413).

$$\forall x \in \mathbb{X} \Rightarrow \exists g \in \mathbb{G} : x = \text{gpm}(g) \quad (1.55)$$

⁶⁵ http://en.wikipedia.org/wiki/Design_pattern_%28computer_science%29 [accessed 2007-08-12]

4. The search space $\mathbb{G}$ should be unbiased in the sense that all phenotypes are represented by the same number of phenotypes. This property allows to efficiently select an unbiased start population, giving the optimizer the chance of reaching all parts of the problem space.

$$\forall x_1, x_2 \in \mathbb{X} \Rightarrow |\{g \in \mathbb{G} : x_1 = \text{gpm}(g)\}| \approx |\{g \in \mathbb{G} : x_2 = \text{gpm}(g)\}| \quad (1.56)$$

5. The genotype-phenotype mapping should always yield *valid* phenotypes. The meaning of valid in this context is that if the problem space $\mathbb{X}$ is the set of all possible trees, only trees should be encoded in the genome. If we use the $\mathbb{R}^3$ as problem space, no vectors with fewer or more elements than three should be produced by the genotype-phenotype mapping. This form of validity does not imply that the individuals are also *correct* solutions in terms of the objective functions.
6. The genotype-phenotype mapping should be simple and bijective.
7. The representations in the search space should possess strong causality (locality), i. e., small changes in the genotype lead to small changes in the phenotype. Optimally, this would mean that:

$$\forall x_1, x_2 \in \mathbb{X}, g \in \mathbb{G} : x_1 = \text{gpm}(g) \wedge x_2 = \text{gpm}(\text{searchOp}(g)) \Rightarrow x_2 \approx x_1 \quad (1.57)$$

Ronald [264] summarizes some further rules [264, 260]:

8. The genotypic representation should be aligned to a set of reproduction operators in a way that good configurations of formae are preserved by the search operations and do not easily get lost during the exploration of the search space.
9. The representations should minimize epistasis (see Section 1.4.5 on page 47 and the 1st rule).
10. The problem should be represented at an appropriate level of abstraction.
11. If a direct mapping between genotypes and phenotypes is not possible, a suitable artificial embryogeny approach should be applied.

Let us now summarize some more conclusions for search spaces based on forma analysis as stated by Radcliffe [257] and Weicker [255].

12. Formae in Genotypic and Phenotypic Space

The optimization algorithms find new elements in the search space $\mathbb{G}$ by applying the search operations $\text{searchOp} \in \text{Op}$. These operations can only create, modify, or combine genotypical formae since they usually have no information about the problem space. Most mathematical models dealing with the propagation of formae like the Building Block Hypothesis and the Schema Theorem⁶⁶ thus focus on the search space and show that highly fit *genotypical* formae will more probably be investigated further than those of low utility. Our goal, however, is to find highly fit formae in the *problem space* $\mathbb{X}$. Such properties can only be created, modified, and combined by the search operations if they correspond to genotypical formae. A good genotype-phenotype mapping should provide this feature.

It furthermore becomes clear that useful separate properties in phenotypic space can only be combined by the search operations properly if they are represented by separate formae in genotypic space too.

13. Compatibility of Formae

Formae of different properties should be compatible. Compatible Formae in phenotypic space should also be compatible in genotypic space. This leads to a low level of epistasis and hence will increase the chance of success of the reproduction operations.

⁶⁶ See Section 3.7 for more information on the Schema Theorem.

14. Inheritance of Formae

The 8th rule mentioned Formae should not get lost during the exploration of the search space. From a good binary search operation like recombination (crossover) in genetic algorithms, we can expect that if its two parameters g_1 and g_2 are members of a forma A , the resulting element will also be an instance of A .

$$\forall g_1, g_2 \in A \subseteq \mathbb{G} \Rightarrow \text{searchOp}(g_1, g_2) \in A \quad (1.58)$$

If we furthermore can assume that all instances of all formae A with minimal precision ($A \in \text{mini}$) of an individual are inherited by at least one parent, the binary reproduction operation is considered as *pure*.

$$\forall g_3 = \text{searchOp}(g_1, g_2) \in \mathbb{G}, \forall A \in \text{mini} : g_3 \in A \Rightarrow g_1 \in A \vee g_2 \in A \quad (1.59)$$

If this is the case, all properties of a genotype g_3 which is a combination of two others g_1, g_2 can be traced back to at least one of its parents. Otherwise, searchOp also performs an implicit unary search step, a mutation in genetic algorithm, for instance. Such properties, although discussed here for binary search operations only, can be extended to arbitrary n -ary operators.

15. Combinations of Formae

If genotypes $g_1, g_2, \dots$ which are instances of different but compatible formae $A_1 \bowtie A_2 \bowtie \dots$ are combined by a binary (or n -ary) search operation, the resulting genotype g should be an instance of both properties, i. e., the combination of compatible formae should be a forma itself.

$$\forall g_1 \in A_1, g_2 \in A_2, \dots \Rightarrow \text{searchOp}(g_1, g_2, \dots) \in A_1 \cap A_2 \cap \dots (\neq \emptyset) \quad (1.60)$$

If this principle holds for many individuals and formae, useful properties can be combined by the optimization step by step, narrowing down the precision of the arising, most interesting formae more and more. This should lead the search to the most promising regions of the search space.

16. Reachability of Formae

The set of available search operations Op should include at least one unary search operation which is able to reach all possible formae. If the binary search operations in Op all are pure, this unary operator is the only one (apart from creation operations) able to introduce new formae which are not yet present in the population. Hence, it should be able to find any given forma.

1.6 General Information

To all the optimization methods that are discussed in this book, you will find such a *General Information* section. Here we outline some of the applications of the respective approach, name the most important conferences, journals, and books as well as link to some online resources.

1.6.1 Areas Of Application

Some example areas of application of global optimization algorithms are:

Application	References
Chemistry, Chemical Engineering	[265, 266, 267]
Biochemistry	[268]
Constraint Satisfaction Problems (CSP)	[7]
Multi-Criteria Decision Making (MCDM)	[20, 18]
Biology	[267]
Engineering, Structural Optimization, and Design	[269, 267, 23, 270, 266, 268, 267, 271]
Economics and Finance	[270, 267, 272]
Parameter Estimation	[268]
Mathematical Problems	[273]
Optics	[274, 275]
Operations Research	[267, 276]
Networking and Communication	[277]
Networking and Communication	Section 23.1 on page 347

This is just a small sample of the possible applications of global optimization algorithms. It has neither some sort of order nor a focus on some specific areas. In the general information sections of the following chapters, you will find many application examples for the algorithm discussed.

1.6.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on global optimization algorithms are:

AAAI: National Conference on Artificial Intelligence

<http://www.aaai.org/Conferences/conferences.php> [accessed 2007-09-06]

History: 2007: Vancouver, British Columbia, Canada, see [278]
 2006: Boston, Massachusetts, USA, see [279]
 2005: Pittsburgh, Pennsylvania, USA, see [280]
 2004: San Jose, California, USA, see [281]
 2002: Edmonton, Alberta, Canada, see [282]
 2000: Austin, Texas, USA, see [283]
 1999: Orlando, Florida, USA, see [284]
 1998: Madison, Wisconsin, USA, see [285]
 1997: Providence, Rhode Island, USA, see [286, 287]
 1996: Portland, Oregon, USA, see [288, 289]
 1994: Seattle, WA, USA, see [290]
 1993: Washington, DC, USA, see [291]
 1992: San Jose, California, USA, see [292]
 1991: Anaheim, California, USA, see [293]
 1990: Boston, Massachusetts, USA, see [294]
 1988: St. Paul, Minnesota, USA, see [295]
 1987: Seattle, WA, USA, see [296]
 1986: Philadelphia, PA, USA, see [297, 298]
 1984: Austin, TX, USA, see [299]

- 1983: Washington, DC, USA, see [300]
- 1982: Pittsburgh, PA, USA, see [301]
- 1980: Stanford University, California, USA, see [302]

HIS: International Conference on Hybrid Intelligent Systems

<http://www.softcomputing.net/hybrid.html> [accessed 2007-09-01]

- History: 2007: Kaiserslautern, Germany, see [303]
 2006: Auckland, New Zealand, see [304]
 2005: Rio de Janeiro, Brazil, see [305]
 2004: Kitakyushu, Japan, see [306]
 2003: Melbourne, Australia, see [307]
 2002: Santiago, Chile, see [308]
 2001: Adelaide, Australia, see [309]

ICNC: International Conference on Advances in Natural Computation

- History: 2007: Haikou, China, see [310, 311, 312, 313, 314]
 2006: Xi'an, China, see [315, 316]
 2005: Changsha, China, see [317, 318, 319]

IAAI: Conference on Innovative Applications of Artificial Intelligence

<http://www.aaai.org/Conferences/IAAI/iaai.php> [accessed 2007-09-06]

- History: 2006: Boston, Massachusetts, USA, see [279]
 2005: Pittsburgh, Pennsylvania, USA, see [280]
 2004: San Jose, California, USA, see [281]
 2003: Acapulco, México, see [320]
 2002: Edmonton, Alberta, Canada, see [282]
 2001: Seattle, Washington, USA, see [321]
 2000: Austin, Texas, USA, see [283]
 1999: Orlando, Florida, USA, see [284]
 1998: Madison, Wisconsin, USA, see [285]
 1997: Providence, Rhode Island, USA, see [286]
 1996: Portland, Oregon, USA, see [288]
 1995: Montreal, Quebec, Canada, see [322]
 1994: Seattle, Washington, USA, see [323]
 1993: Washington, DC, USA, see [324]
 1992: San Jose, California, USA, see [325]
 1991: Anaheim, California, USA, see [326]
 1990: Washington, DC, USA, see [327]
 1989: Stanford University, California, USA, see [328]

KES: Knowledge-Based Intelligent Information & Engineering Systems

- History: 2007: Vietri sul Mare, Italy, see [329, 330, 331]
 2006: Bournemouth, UK, see [332, 333, 334]
 2005: Melbourne, Australia, see [335, 336, 337, 338]
 2004: Wellington, New Zealand, see [339, 340, 341]
 2003: Oxford, UK, see [342, 343]
 2002: Podere d'Ombriano, Crema, Italy, see [344]

- 2001: Osaka and Nara, Japan, see [345]
- 2000: Brighton, UK, see [346, 347]
- 1999: Adelaide, South Australia, see [348]
- 1998: Adelaide, South Australia, see [349, 350, 351]
- 1997: Adelaide, South Australia, see [352, 353]

MCDM: International Conference on Multiple Criteria Decision Making

<http://project.hkikk.fi/MCDM/conf.html> [accessed 2007-09-10]

- History:
- 2008: Auckland, New Zealand, see [354]
 - 2006: Chania, Crete, Greece, see [355]
 - 2004: Whistler, British Columbia, Canada, see [356]
 - 2002: Semmering, Austria, see [357]
 - 2000: Ankara, Turkey, see [358]
 - 1998: Charlottesville, Virginia, USA, see [20]
 - 1997: Cape Town, South Africa, see [359]
 - 1995: Hagen, Germany, see [360]
 - 1994: Coimbra, Portugal, see [361]
 - 1992: Taipei, Taiwan, see [362]
 - 1990: Fairfax, USA, see [363]
 - 1988: Manchester, UK, see [364]
 - 1986: Kyoto, Japan, see [365]
 - 1984: Cleveland, Ohio, USA, see [366]
 - 1982: Mons, Belgium, see [367]
 - 1980: Newark, Delaware, USA, see [368]
 - 1979: Königswinter, Germany, see [369]
 - 1977: Buffalo, New York, USA, see [370]
 - 1975: Jouy-en-Josas, France, see [371]

Mendel: International Conference on Soft Computing

<http://mendel-conference.org/> [accessed 2007-09-09]

- History:
- 2007: Prague, Czech Republic, see [372]
 - 2006: Brno, Czech Republic, see [373]
 - 2005: Brno, Czech Republic, see [374]
 - 2004: Brno, Czech Republic, see [375]
 - 2003: Brno, Czech Republic, see [376]
 - 2002: Brno, Czech Republic, see [377]
 - 2001: Brno, Czech Republic, see [378]
 - 2000: Brno, Czech Republic, see [379]
 - 1999: Brno, Czech Republic, see [380]
 - 1998: Brno, Czech Republic, see [381]
 - 1997: Brno, Czech Republic, see [382]
 - 1996: Brno, Czech Republic, see [383]
 - 1995: Brno, Czech Republic, see [384]

MIC: Metaheuristics International Conference

- History:
- 2007: Montreal, Canada, see [385]

- 2005: Vienna, Austria, see [386]
- 2003: Kyoto, Japan, see [387]
- 2001: Porto, Portugal, see [388]
- 1999: Angra dos Reis, Brazil, see [389]
- 1997: Sophia Antipolis, France, see [390]
- 1995: Breckenridge, Colorado, USA, see [391]

MICAI: Advances in Artificial Intelligence, The Mexican International Conference on Artificial Intelligence

<http://www.micai.org/> [accessed 2008-06-29]

- History: 2007: Aguascalientes, México, see [392]
- 2006: Apizaco, México, see [393, 394]
 - 2005: Monterrey, México, see [395]
 - 2004: Mexico City, México, see [396]
 - 2002: Mérida, Yucatán, México, see [397]
 - 2000: Acapulco, México, see [398]

In the general information sections of the following chapters, you will find many conferences and workshops that deal with the respective algorithms discussed, so this is just a small selection.

1.6.3 Journals

Some journals that deal (at least partially) with global optimization algorithms are:

Journal of Global Optimization, ISSN: 0925-5001 (Print) 1573-2916 (Online), appears monthly, publisher: Springer Netherlands, <http://www.springerlink.com/content/100288/> [accessed 2007-09-20]

The Journal of the Operational Research Society, ISSN: 0160-5682, appears monthly, editor(s): John Wilson, Terry Williams, publisher: Palgrave Macmillan, The OR Society, <http://www.palgrave-journals.com/jors/> [accessed 2007-09-16]

IEEE Transactions on Systems, Man, and Cybernetics (SMC), appears Part A/B: bi-monthly, Part C: quaterly, editor(s): Donald E. Brown (Part A), Diane Cook (Part B), Vladimir Marik (Part C), publisher: IEEE Press, <http://www.ieeesmc.org/> [accessed 2007-09-16]

Journal of Heuristics, ISSN: 1381-1231 (Print), 1572-9397 (Online), appears bi-monthly, publisher: Springer Netherlands, <http://www.springerlink.com/content/102935/> [accessed 2007-09-16]

European Journal of Operational Research (EJOR), ISSN: 0377-2217, appears bi-weekly, editor(s): Roman Slowinski, Jesus Artalejo, Jean-Charles. Billaut, Robert Dyson, Lorenzo Peccati, publisher: North-Holland, Elsevier, http://www.elsevier.com/wps/find/journaldescription.cws_home/505543/description [accessed 2007-09-21]

Computers & Operations Research, ISSN: 0305-0548, appears monthly, editor(s): Stefan Nickel, publisher: Pergamon, Elsevier, http://www.elsevier.com/wps/find/journaldescription.cws_home/300/description [accessed 2007-09-21]

Applied Statistics, ISSN: 0035-9254, editor(s): Gilmour, Skinner, publisher: Blackwell Publishing for the Royal Statistical Society, <http://www.blackwellpublishing.com/journal.asp?ref=0035-9254> [accessed 2007-09-16]

Applied Intelligence, ISSN: 0924-669X (Print), 1573-7497 (Online), appears bi-monthly, publisher: Springer Netherlands, <http://www.springerlink.com/content/100236/> [accessed 2007-09-16]

Artificial Intelligence Review, ISSN: 0269-2821 (Print), 1573-7462 (Online), appears until 2005, publisher: Springer Netherlands, <http://www.springerlink.com/content/100240/> [accessed 2007-09-16]

Journal of Artificial Intelligence Research (JAIR), ISSN: 11076-9757, editor(s): Toby Walsh, <http://www.jair.org/> [accessed 2007-09-16]

Knowledge and Information Systems, ISSN: 0219-1377 (Print), 0219-3116 (Online), appears approx. eight times a year, publisher: Springer London, <http://www.springerlink.com/content/0219-1377> [accessed 2007-09-16] and <http://www.springer.com/west/home/computer/information+systems?SGWID=4-152-70-1136715-0> [accessed 2007-09-16]

SIAM Journal on Optimization (SIOPT), ISSN: 1052-6234 (print) / 1095-7189 (electronic), appears quarterly, editor(s): Nicholas I. M. Gould, publisher: Society for Industrial and Applied Mathematics, <http://www.siam.org/journals/siopt.php> [accessed 2008-06-14]

Applied Soft Computing, ISSN: 1568-4946, appears quarterly, editor(s): R. Roy, publisher: Elsevier B.V., <http://www.sciencedirect.com/science/journal/15684946> [accessed 2008-06-15]

Advanced Engineering Informatics, ISSN: 1474-0346, appears quaterly, editor(s): J.C. Kunz, I.F.C. Smith, T. Tomiyama, publisher: Elsevier B.V., http://www.elsevier.com/wps/find/journaldescription.cws_home/622240/description [accessed 2008-08-01]

Journal of Machine Learning Research (JMLR), ISSN: 1533-7928, 1532-4435, appears 8 times/year, editor(s): Lawrence Saul and Leslie Pack Kaelbling, publisher: Microtome Publishing, <http://jmlr.csail.mit.edu/> [accessed 2008-08-06]

1.6.4 Online Resources

Some general, online available ressources on global optimization algorithms are:

<http://www.mat.univie.ac.at/~neum/glopt.html> [accessed 2007-09-20]

Last update: up-to-date

Description: Arnold Neumaier's global optimization website which includes links, publications, and software.

<http://www.soft-computing.de/> [accessed 2008-05-18]

Last update: up-to-date

Description: Yaochu Jin's site on soft computing including links and conference infos.

<http://web.ift.uib.no/~antonych/glob.html> [accessed 2007-09-20]

Last update: up-to-date

Description: Web site with many links maintained by Gennady A. Ryzhikov.

http://www-optima.amp.i.kyoto-u.ac.jp/member/student/hedar/Hedar_files/TestGO.htm [accessed 2007-11-06]

Last update: up-to-date

Description: A beautiful collection of test problems for global optimization algorithms

1.6.5 Books

Some books about (or including significant information about) global optimization algorithms are:

Pardalos, Thoai, and Horst [8]: *Introduction to Global Optimization*

- Floudas and Pardalos [267]: *Frontiers in Global Optimization*
- Dzemyda, Saltenis, and Zilinskas [270]: *Stochastic and Global Optimization*
- Gandibleux, Sevaux, Sörensen, and T'kindt [399]: *Metaheuristics for Multiobjective Optimization*
- Floudas [268]: *Deterministic Global Optimization: Theory, Methods and Applications*
- Chankong and Haimes [18]: *Multiobjective Decision Making Theory and Methodology*
- Steuer [19]: *Multiple Criteria Optimization: Theory, Computation and Application*
- Haimes, Hall, and Freedman [276]: *Multiobjective Optimization in Water Resource Systems*
- Charnes and Cooper [27]: *Management Models and Industrial Applications of Linear Programming*
- Corne, Dorigo, Glover, Dasgupta, Moscato, Poli, and Price [400]: *New Ideas in Optimisation*
- Jain and Kacprzyk [401]: *New Learning Paradigms in Soft Computing*
- Tiwari, Knowles, Avineri, Dahal, and Roy [402]: *Applications of Soft Computing – Recent Trends*
- Chawdry, Roy, and Pant [271]: *Soft Computing in Engineering Design and Manufacturing*
- Siarry and Michalewicz [403]: *Advances in Metaheuristics for Hard Optimization*
- Onwubolu and Babu [404]: *New Optimization Techniques in Engineering*
- Pardalos and Du [405]: *Handbook of Combinatorial Optimization*
- Reeves [406]: *Modern Heuristic Techniques for Combinatorial Problems*
- Corne, Oates, and Smith [277]: *Telecommunications Optimization: Heuristic and Adaptive Techniques*
- Kontoghiorghes [407]: *Handbook of Parallel Computing and Statistics*
-

Evolutionary Algorithms

*“If your father and his father were fish out of water
You must break the cycle”*

— Illogic [408]

2.1 Introduction

Definition 2.1 (Evolutionary Algorithm). Evolutionary algorithms¹ (EAs) are generic, population-based meta-heuristic optimization algorithms that use biology-inspired mechanisms like mutation, crossover, natural selection, and survival of the fittest. [409, 410, 411]

2.1.1 The Basic Principles from Nature

In 1859, Darwin [6] published his book “On the Origin of Species”² in which he identified the principles of *natural selection* and *survival of the fittest* as driving forces behind the biological evolution. His theory can be condensed into ten observations and deductions [6, 412, 2]:

1. The individuals of a species possess great fertility and produce more offspring than can grow into adulthood.
2. Under the absence of external influences (like natural disasters, human beings, etc.), the population size of a species roughly remains constant.
3. Again, if no external influences occur, the food resources are limited but stable over time.
4. Since the individuals compete for these limited resources, a struggle for survival ensues.
5. Especially in sexual reproducing species, no two individuals are equal.
6. Some of the variations between the individuals will affect their fitness and hence, their ability to survive.
7. Many of these variations are inheritable.
8. Individuals less fit are less likely to reproduce, whereas the fittest individuals will survive and produce offspring more probably.
9. Individuals that survive and reproduce will likely pass on their traits to their offspring.
10. A species will slowly change and adapt more and more to a given environment during this process which may finally even result in new species.

Evolutionary algorithms abstract from this biological process and also introduce a change in semantics by being *goal-driven* [413]. The search space $\mathbb{G}$ in evolutionary algorithms corresponds to the set of all possible DNA strings in nature and its elements $g \in \mathbb{G}$ play the role of the natural genotypes. Therefore, we also often refer to $\mathbb{G}$ as the *genome* and to

¹ http://en.wikipedia.org/wiki/Artificial_evolution [accessed 2007-07-03]

² http://en.wikipedia.org/wiki/The_Origin_of_Species [accessed 2007-07-03]

the elements $g \in \mathbb{G}$ as *genotypes*. Like any creature is an instance of its genotype formed by embryogenesis³, the solution candidates (or *phenotypes*) $x \in \mathbb{X}$ in the problem space $\mathbb{X}$ are instances of genotypes formed by the genotype-phenotype mapping: $x = \text{gpm}(g)$. Their fitness is rated according to objective functions which are subject to optimization and drive the evolution into specific directions.

The advantage of evolutionary algorithms compared to other optimization methods is that they make only few assumptions about the underlying fitness landscape and therefore perform consistently well in many different problem categories. Their basic idea is to copy the process of Darwinian evolution in order to find solutions for hard problems. Let us now take a look on the basic cycle of artificial evolution illustrated in Figure 2.1 and how its single steps correspond with its natural role model.




Fig. 2.1: The basic cycle of evolutionary algorithms.

We can distinguish between single-objective and multi-objective evolutionary algorithms (MOEA), where the latter means that we try to optimize multiple, possibly conflicting criteria. Our following elaborations will be based on MOEAs. The general area of evolutionary computation that deals with multi-objective optimization is called EMOO, evolutionary multi-objective optimization.

Definition 2.2 (MOEA). A multi-objective evolutionary algorithm (MOEA) is able to perform an optimization of multiple criteria on the basis of artificial evolution [414, 14, 415, 12, 13, 15, 91].

All evolutionary algorithms proceed in principle according to the following scheme:

1. Initially, a population Pop of individuals p with a totally random genome $p.g$ is created.
2. The values of the objective functions $f \in F$ are computed for each solution candidate $p.x$ in Pop . This evaluation may incorporate complicated simulations and calculations.
3. With the objective functions, we have determined the utility of the different features of the solution candidates and can now assign a fitness value $v(p.x)$ to each of them. The fitness assignment process incorporates the prevalence comparator function cmp_F which uses the objective values to create a total order amongst the individuals.
4. A subsequent selection process filters out the solution candidates with bad fitness and allows those with good fitness to enter the mating pool with a higher probability. Since fitness is subject to minimization in the context of this book, the lower the $v(p.x)$ -values are, the higher is the (relative) utility of the individual to whom they belong.

³ <http://en.wikipedia.org/wiki/Embryogenesis> [accessed 2008-03-10]

5. In the reproduction phase, offspring is created by varying or combining their genotypes $p.g$ by applying the search operations $\text{searchOp} \in \text{Op}$ which are called *reproduction operations* in the context of EAs. These offspring is then subsequently integrated into the population.
6. If the $\text{terminationCriterion}()$ is met, the evolution stops here. Otherwise, it continues at step 2.

At the beginning of the evolution, there exists no idea what is good or what is bad. Basically, only some random genes $p.x = \text{create}()$ are coupled together as initial population $\text{Pop}(t = 0)$. I think, back in the Eoarchean⁴, the earth age 3.8 billion years ago where most probably the first single-celled life occurred on earth, it was probably the same.

At the beginning of the evolutionary cycle, nature instantiates each genotype $p.g$ in form of a phenotype $p.x = \text{gpm}(p.g)$ – a living organism, for example a fish. The survival of the genes of our fish depend on how good it performs in the ocean ($F(p.x) = ?$), in other words, how fit it is $v(p.x)$. This fitness however is not only determined by one single feature of the phenotype like its size ($= f_1$). Although a bigger fish will have better chances to survive, size alone does not help if it is too slow to catch any prey ($= f_2$). Also its energy consumption f_3 should be low so it does not need to eat all the time. Other factors influencing the fitness positively are formae like sharp teeth f_4 and colors that blend into the environment f_5 so it cannot be seen too easily by sharks. But, uh, wait a second, if its camouflage is too good, how will it find potential mating partners ($f_6 \approx f_5$)? And if it is big, it will also have a higher energy consumption $f_1 \approx f_3$. So there may be conflicts between the desired properties.

To sum it up, we could consider the life of the fish as the evaluation process of its genotype in an environment where good qualities in one aspect can turn out as drawbacks in other perspectives. In multi-objective evolutionary algorithms, this is exactly the same which we tried to demonstrate by annotating our fish-story with the symbols from global optimization theory. For each problem that we want to solve, we can specify multiple so-called objective functions $f \in F$. An objective function f represents one feature that we are interested in. Let us assume that we want to evolve a car (a pretty weird assumption, but let's stick with it). The genotype $p.g \in \mathbb{G}$ would be the construction plan and the phenotype $p.x \in \mathbb{X}$ the real car, or at least a simulation of it. One objective function f_a would definitely be *safety*. For the sake of our children and their children, the car should also be *environment-friendly*⁵, so that's our second objective function f_b . Furthermore, a cheap price f_c , fast speed f_d , and a cool design f_e would be good. That makes five objective functions from which for example the second and the fourth are contradictory ($f_b \approx f_d$).

After the fish genome is instantiated, nature “knows” about its phenotypic properties. Fitness, however, is always relative; it depends on your environment. I, for example, may be considered as a fit man in my department (computer science). If took a stroll to the department of sports science, that statement will probably not hold anymore. The same goes for the fish, its fitness depends on the other fish in the population. If one fish $p_1.x$ can beat another one $p_2.x$ in all categories, i. e., is bigger, stronger, smarter, and so on, we can clearly consider it as fitter ($p_1.x \succ p_2.x \Rightarrow \text{cmp}_F(p_1.x, p_2.x) < 0$) since it will have a better chance to survive. This relation is transitive but only forms a partial order since a fish that is strong but not very clever and a fish that is clever but not strong maybe have the same probability to reproduce and hence, are not directly comparable⁶. Well, Ok, we cannot decide if a weak fish $p_3.x$ with a clever behavioral pattern is worse or better than a really strong but less cunning one $p_4.x$ ($\text{cmp}_F(p_3.x, p_4.x) = 0$). Both traits are furthered in the evolutionary process and maybe, one fish of the first kind will sometimes mate with one of the latter and produce an offspring which is both, intelligent and sporty⁷.

⁴ <http://en.wikipedia.org/wiki/Eoarchean> [accessed 2007-07-03]

⁵ Get my point?

⁶ Which is a very comforting thought for all computer scientists.

⁷ I wonder if the girls in the sports department are open to this kind of argumentation?

Multi-objective evolutionary algorithms basically apply the same principles in their fitness assignment process “assignFitness”. One of the most popular methods for computing the fitness here is called *Pareto ranking*⁸. It does exactly what we’ve just discussed: It first chooses the individuals that are beaten by no one (we call this non-dominated set) and assigns a good (scalar) fitness value $v(p_1.x)$ to them. Then it looks at the rest of the population and picks those ($P \subset Pop$) which are not beaten by the remaining individuals and gives them a slightly worse fitness value $v(p.x) > v(p_1.x) \forall p \in P$ – and so on, until all solution candidates have received one scalar fitness.

Now, how fit a fish is does not necessarily determine directly if it can produce offspring. An intelligent fish may be eaten by a shark and a strong one can die from disease. The fitness is only some sort of probability of reproduction. The process of selection is always stochastic, without guarantees – even a fish that is small, slow, and lacks any sophisticated behavior might survive and could produce even more offspring than a highly fit one.

The selection in evolutionary algorithms works in exactly the same way – it uses a selection algorithm “select” in order to pick the fittest individuals and places them into the mating pool *Mate*. The oldest selection scheme is called *Roulette wheel*⁹. In the original version of this algorithm (intended for fitness maximization), the chance of an individual p to reproduce was proportional to its fitness $v(p.x)$.

Last but not least, there is the reproduction phase. Fish reproduce sexually. Whenever a female fish and a male fish mate, their genes will be recombined by crossover. Furthermore, mutations may take place which. Most often, they affect the characteristics of resulting larva only slightly [416]. Since fit fish produce offspring with higher probability, there is a good chance that the next generation will contain at least some individuals that have combined good traits from their parents and perform even better than them.

In evolutionary algorithms, we do not have such a thing as “gender”. Each individual from the mating pool can potentially be recombined with every other one. In the car example, this means that we would modify the construction plans by copying the engine of one car and placing it into the car body of another one. Also, we could alter some features like the shape of the headlights randomly. This way, we receive new construction plans for new cars. Our chance that an *environment-friendly* engine inside a *cool-looking* car will result in a car that is more likely to be bought by the customer is good. If we iteratively perform the reproduction process “reproducePop” time and again, there is a high probability that the solutions finally found will be close to optimal.

From the Viewpoint of Formae Analysis

Let us review our introductory fish example in terms of forma analysis. Fish can, for instance, be characterized by the properties “clever” and “strong”. Crudely simplified, both properties may be **true** or **false** for a single individual and hence define two formae each. A third property can be the color, for which many different possible variations exist. Some of them may be good in terms of camouflage, others maybe good in terms of finding mating partners. Now a fish can be clever and strong at the same time, as well as weak and green. Here, a living fish allows nature to evaluate the utility of at least three different formae.

This fact has first been stated by Holland [253] for genetic algorithms and is termed *implicit parallelism* (or *intrinsic parallelism*). Since then, it has been studied by many different researchers [417, 418, 419, 420]. If the search space and the genotype-phenotype mapping are properly designed, the implicit parallelism in conjunction with the crossover/recombination operations is one of the reasons why evolutionary algorithms are such a successful class of optimization algorithms.

⁸ Pareto comparisons are discussed in Section 1.2.4 on page 13 and elaborations on Pareto ranking can be found in Section 2.3.3.

⁹ The roulette wheel selection algorithm will be introduced in Section 2.4.3 on page 95.

Does the natural Paragon Fit?

At this point it should be mentioned that the direct reference to Darwinian evolution in evolutionary algorithms is somehow controversial. Paterson [44], for example, points out that “neither GAs [genetic algorithms] nor GP [Genetic Programming] are concerned with the evolution of new species, nor do they use natural selection.” On the other hand, nobody would claim that the idea of selection has not been borrowed from nature although many additions and modifications have been introduced in favor for better algorithmic performance. The second argument concerning the development of different species depends on definition: According to Wikipedia [2], a species is a class of organisms which are very similar in many aspects such as appearance, physiology, and genetics. In principle, there is some elbowroom for us and we may indeed consider even different solutions to a single problem in evolutionary algorithms as members of a different species – especially if the binary search operation crossover/recombination applied to their genomes cannot produce another valid solution candidate.

Another interesting difference was pointed out by Sharpe [421] who states that natural evolution “only proceed[s] sufficiently fast to ensure survival” whereas evolutionary algorithms used for engineering need to be fast in order to be feasible and to compete with other problem solving techniques.

Furthermore, although the concept of fitness¹⁰ in nature is controversial [46], it is often considered as an *a posteriori measurement*. It then defines the ratio of the numbers of occurrences of a genotype in a population after and before selection or the number of offspring an individual has in relation to the number of offspring of another individual. In evolutionary algorithms, fitness is an *a priori quantity* denoting a value that determines the expected number of instances of a genotype that should survive the selection process. However, one could conclude that biological fitness is just an approximation of the *a priori* quantity arisen due to the hardness (if not impossibility) of directly measuring it.

My personal opinion (which may as well be wrong) is that the citation of Darwin here is well motivated since there are close parallels between Darwinian evolution and evolutionary algorithms. Nevertheless, natural and artificial evolution are still two different things and phenomena observed in either of the two do not necessarily carry over to the other.

The Basic Evolutionary Algorithm Scheme

After this informal outline about the artificial evolution and how we can use it as an optimization method, let us now specify the basic scheme common to all evolutionary algorithms. In principle, all EAs are variations and extensions of the basic approach “simpleEA” defined Algorithm 2.1, a cycle of evaluation, selection, and reproduction repeated in each iteration t . Algorithm 2.1 relies on functions and prototypes that we will introduce step by step.

- The function “createPop(ps)”, which will be introduced as Algorithm 2.17 in Section 2.5 on page 108, produces an initial, randomized population consisting of ps individuals in the first iteration $t = 0$.
- The termination criterion “terminationCriterion()” checks whether the evolutionary algorithm should terminate or continue its work, see Section 1.3.4 on page 35.
- Most evolutionary algorithms assign a scalar fitness $v(p.x)$ to each individual p by comparing its vector of objective values $F(p.x)$ to other individuals in the population Pop . The function v is built by a fitness assignment process “assignFitness”, which we will discuss in Section 2.3 on page 82 in more detail. During this procedure, the genotype-phenotype mapping is implicitly carried out as well as simulations needed to compute the objective functions $f \in F$.
- A selection algorithm “select” (see Section 2.4 on page 92) then chooses ps interesting individuals from the population Pop and inserts them into the mating pool $Mate$.

¹⁰ [http://en.wikipedia.org/wiki/Fitness_\(biology\)](http://en.wikipedia.org/wiki/Fitness_(biology)) [accessed 2008-08-10]

Algorithm 2.1: $X^* \leftarrow \text{simpleEA}(\text{cmp}_F, ps)$

Input: cmp_F : the comparator function which allows us to compare the utility of two solution candidates
Input: ps : the population size
Data: t : the generation counter
Data: Pop : the population
Data: $Mate$: the mating pool
Data: v : the fitness function resulting from the fitness assigning process
Output: X^* : the set of the best elements found

```

1 begin
2    $t \leftarrow 0$ 
3    $Pop \leftarrow \text{createPop}(ps)$ 
4   while  $\text{terminationCriterion}()$  do
5      $v \leftarrow \text{assignFitness}(Pop, \text{cmp}_F)$ 
6      $Mate \leftarrow \text{select}(Pop, v, ps)$ 
7      $t \leftarrow t + 1$ 
8      $Pop \leftarrow \text{reproducePop}(Mate)$ 
9   return  $\text{extractPhenotypes}(\text{extractOptimalSet}(Pop))$ 
10 end

```

- With “reproducePop”, a new population is generated from the individuals inside the mating pool using mutation and/or recombination. More information on reproduction can be found in Section 2.5 on page 108 and in Definition 2.13.
- The functions “extractOptimalSet” and “extractPhenotypes” which you can find introduced in Definition 18.2 on page 264 and Equation 18.1 on page 263 are used to extract all the non-prevalled individuals p^* from the final population and to return their corresponding phenotypes $p^*.x$ only.

2.1.2 Classification of Evolutionary Algorithms

The Family of Evolutionary Algorithms

The family of evolutionary algorithms encompasses five members, as illustrated in Figure 2.2. We will only enumerate them here in short. In depth discussions will follow in the next chapters.

- **Genetic algorithms** (GAs) are introduced in Chapter 3 on page 113. GAs subsume all evolutionary algorithms which have bit strings as search space $\mathbb{G}$.
- The set of evolutionary algorithms which explore the space of real vectors $\mathbb{X} \subseteq \mathbb{R}^n$ is called **Evolution Strategies** (ES, see Chapter 5 on page 191).
- For **Genetic Programming** (GP), which will be elaborated on in Chapter 4 on page 127, we can provide two definitions: On one hand, GP includes all evolutionary algorithms that grow programs, algorithms, and these alike. On the other hand, also all EAs that evolve tree-shaped individuals are instances of Genetic Programming.
- **Learning Classifier Systems** (LCS), discussed in Chapter 7 on page 197, are online learning approaches that assign output values to given input values. They internally use a genetic algorithm to find new rules for this mapping.
- **Evolutionary programming** (EP, see Chapter 6 on page 195) is an evolutionary approach that treats the instances of the genome as different species rather than as individuals. Over the decades, it has more or less merged into Genetic Programming and the other evolutionary algorithms.

The early research [422] in genetic algorithms (see Section 3.1 on page 113), Genetic Programming (see Section 4.1.1 on page 127), and evolutionary programming (see Section 6.1




Fig. 2.2: The family of evolutionary algorithms.

on page 195) date back to the 1950s and 60s. Besides the pioneering work listed in these sections, at least other important early contribution should not go unmentioned here: The Evolutionary Operation (EVOP) approach introduced by Box [423], Box and Draper [424] in the late 1950s. The idea of EVOP was to apply a continuous and systematic scheme of small changes in the control variables of a process. The effects of these modifications are evaluated and the process is slowly shifted into the direction of improvement. This idea was never realized as a computer algorithm, but Spendley et al. [425] used it as basis for their simplex method which then served as progenitor of the downhill simplex algorithm¹¹ of Nelder and Mead [426]. [422, 427] Satterthwaite's REVOP [428, 429], a randomized Evolutionary Operation approach, however, was rejected at this time [422].

We now have classified different evolutionary algorithms according to their semantics, in other words, corresponding to their special search and problem spaces. All five major approaches can be realized with the basic scheme defined in Algorithm 2.1. To this simple structure, there exist many general improvements and extensions. Since these normally do not concern the search or problem spaces, they also can be applied to all members of the EA family alike. In the further text of this chapter, we will discuss the major components of many of today's most efficient evolutionary algorithms [430]. The distinctive features of these EAs are:

- The population size or the number of populations used.
- The method of selecting the individuals for reproduction.
- The way the offspring is included into the population(s).

Populations in Evolutionary Algorithms

There exist various way in which an evolutionary algorithm can process its population. Especially interesting is how the population $Pop(t+1)$ of the next iteration is formed as a combination of the current one $Pop(t)$ and its offspring. If it only contains this offspring, we speak of *extinctive selection* [431, 432]. Extinctive selection can be compared with ecosystems

¹¹ We discuss Nelder and Mead [426]'s downhill simplex optimization method in Chapter 11 on page 223.

of small protozoa¹² which reproduce in a fissiparous¹³ manner. In this case, of course, the elders will not be present in the next generation. Other comparisons can partly be drawn to the sexual reproducing to octopi, where the female dies after protecting the eggs until the larvae hatch, or to the black widow spider where the female devours the male after the insemination. Especially in the area of genetic algorithms, extinctive strategies are also known as *generational algorithms*.

Definition 2.3 (Generational). In evolutionary algorithms that are *generational* [433], the next generation will only contain the offspring of the current one and no parent individuals will be preserved.

Extinctive evolutionary algorithms can further be divided into *left* and *right* selection [434]. In left extinctive selections, the best individuals are not allowed to reproduce in order to prevent premature convergence of the optimization process. Conversely, the worst individuals are not permitted to breed in right extinctive selection schemes in order to reduce the selective pressure since they would otherwise scatter the fitness too much.

In algorithms that apply a *preservative selection* scheme, the population is a combination of the next population and the offspring [435, 436, 437, 413]. The biological metaphor for such algorithms is that the lifespan of many organisms exceeds a single generation. Hence, parent and child individuals compete with each other for survival.

For Evolution Strategy which you can find discussed in Chapter 5 on page 191, there exists a notation which also can be used describe the generation transition in evolutionary algorithms in general [438, 439, 440, 435].

- λ denotes the number of offspring created and
- μ is the number of parent individuals.

Extinctive selection patterns are denoted as (μ, λ) -strategies and will create $\lambda \geq \mu$ child individuals from the μ available genotypes. From these, they only keep the μ best solution candidates and discard the μ parents as well as the $\lambda - \mu$ worst children.

In $(\mu + \lambda)$ -strategy, again λ children are generated from μ parents, often with $\lambda > \mu$. Then, the parent and offspring populations are united (to a population of the size $\lambda + \mu$) and from this unison, only the μ best individuals will “survive”. $(\mu + \lambda)$ -strategies are thus preservative.

Steady-state evolutionary algorithms [441, 442, 443, 444, 445, 446], abbreviated by SSEA, are preservative evolutionary algorithms with values of λ that are relatively low in comparison with μ . Usually, λ is chosen in a way that a binary search operator crossover is applied exactly once per generation. Although steady-state evolutionary algorithms are often observed to produce better results than generational EAs. Chafekar et al. [444], for example, introduce steady-state evolutionary algorithms that are able to outperform generational NSGA-II (which you can find summarized in ?? on page ??) for some difficult problems. In experiments of Jones and Soule [447] (primarily focused on other issues), steady-state algorithms showed better convergence behavior in a multi-modal landscape. Similar results have been reported by Chevreux [448] in the context of molecule design optimization. Different generational selection methods have been compared to the steady-state GENITOR approach by Goldberg and Deb [449]. On the other hand, with steady-state approaches, we run also the risk of premature convergence.

Even in preservative strategies, it is not granted that the best individuals will always survive. In principle, a $(\mu + \lambda)$ strategy can also mean that from $\mu + \lambda$ individuals, μ are chosen with a certain selection algorithm. Most are randomized, and even if such methods pick the best solution candidates with the highest probabilities, they may also select worse individuals. At this point, it is maybe interesting to mention that the idea that larger populations will always lead to better optimization results does not necessarily always hold, as shown by van Nimwegen and Crutchfield [172].

¹² <http://en.wikipedia.org/wiki/Protozoa> [accessed 2008-03-12]

¹³ http://en.wikipedia.org/wiki/Binary_fission [accessed 2008-03-12]

Definition 2.4 (Elitism). An elitist evolutionary algorithm [254, 450, 414] ensures that at least one copy of the best individual(s) of the current generation is propagated on to the next generation.

The main advantage of elitism is that its convergence is guaranteed, meaning that once the global optimum has been discovered, the evolutionary algorithm converges to that optimum. On the other hand, the risk of converging to a local optimum is also higher. Elitism is an additional feature of global optimization algorithms – a special type of preservative strategy – which is often realized by using a secondary population only containing the non-prevalled individuals. This population is updated at the end of each iteration. Such an archive-based elitism can be combined with both, generational and preservative strategies. Algorithm 2.2 specifies the basic scheme of elitist evolutionary algorithms.

Algorithm 2.2: $X^* \leftarrow \text{elitistEA}(\text{cmp}_F, ps, a)$

Input: cmp_F : the comparator function which allows us to compare the utility of two solution candidates
Input: ps : the population size
Input: as : the archive size
Data: t : the generation counter
Data: Pop : the population
Data: $Mate$: the mating pool
Data: Arc : the archive with the best individuals found so far
Data: v : the fitness function resulting from the fitness assigning process
Output: X^* : the set of best solution candidates discovered

```

1 begin
2    $t \leftarrow 0$ 
3    $Arc \leftarrow \emptyset$ 
4    $Pop \leftarrow \text{createPop}(ps)$ 
5   while  $\text{terminationCriterion}()$  do
6      $Arc \leftarrow \text{updateOptimalSetN}(Arc, Pop)$ 
7      $Arc \leftarrow \text{pruneOptimalSet}(Arc, as)$ 
8      $v \leftarrow \text{assignFitness}(Pop, Arc, \text{cmp}_F)$ 
9      $Mate \leftarrow \text{select}(Pop, Arc, v, ps)$ 
10     $t \leftarrow t + 1$ 
11     $Pop \leftarrow \text{reproducePop}(Mate)$ 
12  return  $\text{extractPhenotypes}(\text{extractOptimalSet}(Pop \cup Arc))$ 
13 end

```

Let us now outline the new methods and changes introduced in Algorithm 2.2 in short.

- The archive Arc is the set of best individuals found by the algorithm. Initially, it is the empty set $\emptyset$. Subsequently, it is updated with the function “updateOptimalSetN” which inserts new, unprevailed elements from the population into it and also removes individuals from the archive which are superseded by those new optima. Algorithms that realize such updating are defined in Section 18.1 on page 263.
- If the optimal set becomes too large – it might theoretically contain uncountable many individuals – “pruneOptimalSet” reduces it to a proper size, employing techniques like clustering in order to preserve the element diversity. More about pruning can be found in Section 18.3 on page 265.
- You should also notice that both, the fitness assignment and selection processes, of elitist evolutionary algorithms may take the archive as additional parameter. In principle, such archive-based algorithms can also be used in non-elitist evolutionary algorithms by simply replacing the parameter Arc with $\emptyset$.

2.1.3 Configuration Parameters of evolutionary algorithms

Figure 2.3 illustrates the basic configuration parameters of evolutionary algorithms. The performance and success of an evolutionary optimization approach applied to a problem given by a set of objective functions F and a problem space $\mathbb{X}$ is defined by




Fig. 2.3: The configuration parameters of evolutionary algorithms.

1. its basic parameter settings like the population size ps or the crossover and mutation rates,
2. whether it uses an archive Arc of the best individuals found and, if so, which pruning technology is used to prevent it from overflowing,
3. the fitness assignment process “assignFitness” and the selection algorithm “select”,
4. the choice of the search space $\mathbb{G}$ and the search operations Op ,
5. and the genotype-phenotype mapping connecting the search Space and the problem space.

In Section 19.1, we go more into detail on how to state the configuration of an optimization algorithm in order to fully describe experiments and to make them reproducible.

2.2 General Information

2.2.1 Areas Of Application

Some example areas of application of evolutionary algorithms are:

Application	References
Function Optimization	[451, 452]

Multi-Objective Optimization	[33, 15, 430, 453, 31, 13]
Combinatorial Optimization	[89, 437, 454]
Engineering, Structural Optimization, and Design	[455, 456, 457]
Constraint Satisfaction Problems (CSP)	[413, 453, 15, 31]
Economics and Finance	[458, 60, 459, 460, 461]
Biology	[69, 462]
Data Mining and Data Analysis	[463, 464, 465]
Mathematical Problems	[466]
Electrical Engineering and Circuit Design	[467, 69]
Chemistry, Chemical Engineering	[468, 469, 448]
Scheduling	[470, 471, 472, 473, 474]
Robotics	[475]
Image Processing	[476, 477]
Networking and Communication	[478, 479, 480, 481, 482, 483, 484, 485, 486] Section 23.1 on page 347
Medicine	[487, 488]
Ressource Minimization, Environment Surveillance/Protection	[206]
Military and Defense	[489]
Evolving Behaviors, e.g., for Agents or Game Players	[490]

For more information see also the application sections of the different members of the evolutionary algorithm family: genetic algorithms in Section 3.2.1 on page 114, Genetic Programming in Section 4.2.1 on page 130, Evolution Strategy in Section 5.2.1 on page 191, evolutionary programming in Section 6.2.1 on page 195, and Learning Classifier Systems in Section 7.2.1 on page 197.

2.2.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on evolutionary algorithms are:

<i>BIOMA</i> : International Conference on Bioinspired Optimization Methods and their Applications
http://bioma.ijs.si/ [accessed 2007-06-30]
History: 2006: Ljubljana, Slovenia, see [491]
2004: Ljubljana, Slovenia, see [492]
<i>CEC</i> : Congress on Evolutionary Computation
http://ieeexplore.ieee.org/servlet/opac?punumber=7875 [accessed 2007-09-05]
History: 2008: Hong Kong, China, see [493]
2007: Singapore, see [494]
2006: Vancouver, BC, Canada, see [495]
2005: Edinburgh, Scotland, UK, see [496]
2004: Portland, Oregon, USA, see [497]
2003: Canberra, Australia, see [498]
2002: Honolulu, HI, USA, see [499]
2001: Seoul, Korea, see [500]
2000: La Jolla, California, USA, see [501]
1999: Washington D.C., USA, see [502]

- 1998: Anchorage, Alaska, USA, see [503]
- 1997: Indianapolis, IN, USA, see [504]
- 1996: Nagoya, Japan, see [505]
- 1995: Perth, Australia, see [506]
- 1994: Orlando, Florida, USA, see [507]

Dagstuhl Seminar: Practical Approaches to Multi-Objective Optimization

- History: 2006: Dagstuhl, Germany, see [508]
- 2004: Dagstuhl, Germany, see [509]

EA/AE: Conference on Artificial Evolution (Evolution Artificielle)

- History: 2007: Tours, France, see [510]
- 2005: Lille, France, see [511]
- 2003: Marseilles, France, see [512]
- 2001: Le Creusot, France, see [513]
- 1999: Dunkerque, France, see [514]
- 1997: Nîmes, France, see [515]
- 1995: Brest, France, see [516]
- 1994: Toulouse, France, see [517]

EMO: International Conference on Evolutionary Multi-Criterion Optimization

- History: 2007: Matsushima/Sendai, Japan, see [518]
- 2005: Guanajuato, México, see [519]
- 2003: Faro, Portugal, see [520]
- 2001: Zurich, Switzerland, see [521]

EUROGEN: Evolutionary Methods for Design Optimization and Control with Applications to Industrial Problems

- History: 2007: Jyväskylä, Finland, see [522]
- 2005: Munich, Germany, see [523]
- 2003: Barcelona, Spain, see [524]
- 2001: Athens, Greece, see [525]
- 1999: Jyväskylä, Finland, see [526]
- 1997: Trieste, Italy, see [527]
- 1995: Las Palmas de Gran Canaria, Spain, see [528]

EvoCOP: European Conference on Evolutionary Computation in Combinatorial Optimization

- <http://www.evostar.org/> [accessed 2007-09-05]
- Co-located with EvoWorkshops and EuroGP.
- History: 2007: Valencia, Spain, see [529]
- 2006: Budapest, Hungary, see [530]
- 2005: Lausanne, Switzerland, see [531]
- 2004: Coimbra, Portugal, see [532]
- 2003: Essex, UK, see [533]
- 2002: Kinsale, Ireland, see [534]
- 2001: Lake Como, Milan, Italy, see [535]

EvoWorkshops: Applications of Evolutionary Computing: EvoCoMnet, EvoFIN, EvoIASP, EvoINTERACTION, EvoMUSART, EvoPhD, EvoSTOC and EvoTransLog

- <http://www.evostar.org/> [accessed 2007-08-05]
- Co-located with EvoCOP and EuroGP.

History: 2008: Naples, Italy, see [536]
 2007: Valencia, Spain, see [537]
 2006: Budapest, Hungary, see [538]
 2005: Lausanne, Switzerland, see [539]
 2004: Coimbra, Portugal, see [540]
 2003: Essex, UK, see [533]
 2002: Kinsale, Ireland, see [534]
 2001: Lake Como, Milan, Italy, see [535]
 2000: Edinburgh, Scotland, UK, see [541]
 1999: Göteborg, Sweden, see [542]
 1998: Paris, France, see [543]

FEA: International Workshop on Frontiers in Evolutionary Algorithms

Was part of the Joint Conference on Information Science

History: 2005: Salt Lake City, Utah, USA, see [544]
 2003: Cary, North Carolina, USA, see [545]
 2002: Research Triangle Park, North Carolina, USA, see [546]
 2000: Atlantic City, NJ, USA, see [547]
 1998: Research Triangle Park, North Carolina, USA, see [548]
 1997: Research Triangle Park, North Carolina, USA, see [549]

GECCO: Genetic and Evolutionary Computation Conference

<http://www.sigevo.org/> [accessed 2007-08-30]

A recombination of the Annual Genetic Programming Conference (GP, see Section 4.2.2 on page 131) and the International Conference on Genetic Algorithms (ICGA, see Section 3.2.2 on page 115), also “contains” the International Workshop on Learning Classifier Systems (IWLCS, see Section 7.2.2 on page 198).

History: 2008: Atlanta, Georgia, USA, see [550, 461, 489, 490]
 2007: London, England, see [551, 552]
 2006: Seattle, Washington, USA, see [553]
 2005: Washington, D.C., USA, see [554, 555, 556, 557]
 2004: Seattle, Washington, USA, see [558, 559, 560]
 2003: Chicago, Illinois, USA, see [561, 562]
 2002: New York, USA, see [563, 564, 565, 566, 567]
 2001: San Francisco, California, USA, see [568, 569]
 2000: Las Vegas, Nevada, USA, see [570, 571]
 1999: Orlando, Florida, USA, see [572, 573, 478]

ICANNGA: International Conference on Adaptive and Natural Computing Algorithms

before 2005: International Conference on Artificial Neural Nets and Genetic Algorithms

History: 2007: Warsaw, Poland, see [574, 575]
 2005: Coimbra, Portugal, see [576]
 2003: Roanne, France, see [577]
 2001: Prague, Czech Republic, see [578]
 1999: Portoroz, Slovenia, see [579]
 1997: Norwich, England, see [580]
 1995: Alès, France, see [581]
 1993: Innsbruck, Austria, see [582]

ICNC: International Conference on Advances in Natural Computation

see Section 1.6.2 on page 62

Mendel: International Conference on Soft Computing

see Section 1.6.2 on page 63

PPSN: International Conference on Parallel Problem Solving from Nature

<http://ls11-www.informatik.uni-dortmund.de/PPSN/> [accessed 2007-09-05]

History: 2008: Dortmund, Germany, see [583]
 2006: Reykjavik, Iceland, see [584]
 2004: Birmingham, UK, see [585]
 2002: Granada, Spain, see [586]
 2000: Paris, France, see [587]
 1998: Amsterdam, The Netherlands, see [588]
 1996: Berlin, Germany, see [589]
 1994: Jerusalem, Israel, see [590]
 1992: Brussels, Belgium, see [591]
 1990: Dortmund, Germany, see [592]

2.2.3 Journals

Some journals that deal (at least partially) with evolutionary algorithms are:

Evolutionary Computation, ISSN: 1063-6560, appears quarterly, editor(s): Marc Schoenauer, publisher: MIT Press, <http://www.mitpressjournals.org/loi/evco> [accessed 2007-09-16]

IEEE Transactions on Evolutionary Computation, ISSN: 1089-778X, appears bi-monthly, editor(s): Xin Yao, publisher: IEEE Computational Intelligence Society, <http://iee-cis.org/pubs/tec/> [accessed 2007-09-16]

Biological Cybernetics, ISSN: 0340-1200 (Print), 1432-0770 (Online), appears bi-monthly, publisher: Springer Berlin/Heidelberg, <http://www.springerlink.com/content/100465/> [accessed 2007-09-16]

Complex Systems, ISSN: 0891-2513, appears quarterly, editor(s): Stephen Wolfram, publisher: Complex Systems Publications, Inc., <http://www.complex-systems.com/> [accessed 2007-09-16]

Journal of Artificial Intelligence Research (JAIR) (see Section 1.6.3 on page 65)

New Mathematics and Natural Computation (NMNC), ISSN: 1793-0057, appears three times a year, editor(s): Paul P. Wang, publisher: World Scientific, <http://www.worldscinet.com/nmnc/> [accessed 2007-09-19]

The Journal of the Operational Research Society (see Section 1.6.3 on page 64)

2.2.4 Online Resources

Some general, online available resources on evolutionary algorithms are:

<http://www.lania.mx/~ccoello/EMOO/> [accessed 2007-09-20]

Last update: up-to-date

Description: EMOO Web page – Dr. Coello Coello’s giant bibliography and paper repository for evolutionary multi-objective optimization.

http://www-isf.maschinenbau.uni-dortmund.de/links/ci_links.html [accessed 2007-10-14]

Last update: up-to-date

Description: Computational Intelligence (CI)-related links and literature, maintained by Jörn Mehnen

http://www.aip.de/~ast/EvolCompFAQ/ [accessed 2007-09-16]
Last update: 2001-04-01
Description: Frequently Asked Questions of the comp.ai.genetic group by Heitkötter and Beasley [1].
http://nknucc.nknu.edu.tw/~hcwu/pdf/evolec.pdf [accessed 2007-09-16]
Last update: 2005-02-19
Description: Lecture Notes on Evolutionary Computation by Wu [434]
http://ls11-www.cs.uni-dortmund.de/people/beyer/EA-glossary/ [accessed 2008-04-10]
Last update: 2002-02-25
Description: Online glossary on terms and definitions in evolutionary algorithms by Beyer et al. [593]
http://www.illegal.uiuc.edu/web/ [accessed 2008-05-17]
Last update: up-to-date
Description: The Illinois Genetic Algorithms Laboratory (IlligAL)
http://www.peterindia.net/Algorithms.html [accessed 2008-05-17]
Last update: up-to-date
Description: A large collection of links about evolutionary algorithms, Genetic Programming, genetic algorithms, etc.
http://www.fmi.uni-stuttgart.de/fk/evolalg/ [accessed 2008-05-17]
Last update: 2003-07-08
Description: The Evolutionary Computation repository of the University of Stuttgart.
http://dis.ijs.si/filipic/ec/ [accessed 2008-05-18]
Last update: 2007-11-09
Description: The Evolutionary Computation repository of the Jožf Stefan Institute in Slovenia
http://www.red3d.com/cwr/evolve.html [accessed 2008-05-18]
Last update: 2002-07-27
Description: Evolutionary Computation and its application to art and design by Craig Reynolds
http://surf.de.uu.net/encore/ [accessed 2008-05-18]
Last update: 2004-08-26
Description: ENCORE, the electronic appendix to The Hitch-Hiker's Guide to Evolutionary Computation, see [1]
http://www-isf.maschinenbau.uni-dortmund.de/links/ci_links.html [accessed 2008-05-18]
Last update: 2006-09-13
Description: A collection of links to computational intelligence / EAs

2.2.5 Books

Some books about (or including significant information about) evolutionary algorithms are:

Bäck [409]: <i>Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms</i>
Bäck, Fogel, and Michalewicz [410]: <i>Handbook of Evolutionary Computation</i>
Fogel [594]: <i>Evolutionary Computation: The Fossil Record</i>
Deb [13]: <i>Multi-Objective Optimization Using Evolutionary Algorithms</i>
Coello Coello and Lamont [595]: <i>Applications of Multi-Objective Evolutionary Algorithms</i>
Eiben and Smith [56]: <i>Introduction to Evolutionary Computing</i>
Dumitrescu, Lazzerini, Jain, and Dumitrescu [596]: <i>Evolutionary Computation</i>

- Bäck, Fogel, and Michalewicz [597]: *Evolutionary Computation 1: Basic Algorithms and Operators*
- Bäck, Fogel, and Michalewicz [598]: *Evolutionary Computation 2: Advanced Algorithms and Operators*
- Bentley [599]: *Evolutionary Design by Computers*
- De Jong [600]: *Evolutionary Computation: A Unified Approach*
- Weicker [255]: *Evolutionäre Algorithmen*
- Gerdes, Klawonn, and Kruse [601]: *Evolutionäre Algorithmen*
- Nissen [602]: *Einführung in evolutionäre Algorithmen: Optimierung nach dem Vorbild der Evolution*
- Yao [603]: *Evolutionary Computation: Theory and Applications*
- Yu, Davis, Baydar, and Roy [604]: *Evolutionary Computation in Practice*
- Yang, Ong, and Jin [228]: *Evolutionary Computation in Dynamic and Uncertain Environments*
- Morrison [234]: *Designing Evolutionary Algorithms for Dynamic Environments*
- Branke [232]: *Evolutionary Optimization in Dynamic Environments*
- Nedjah, Alba, and Mourelle [431]: *Parallel Evolutionary Computations*
- Rothlauf [115]: *Representations for Genetic and Evolutionary Algorithms*
- Banzhaf and Eckman [605]: *Evolution and Biocomputation – Computational Models of Evolution*
- Fogel and Corne [462]: *Evolutionary Computation in Bioinformatics*
- Johnston [468]: *Applications of Evolutionary Computation in Chemistry*
- Clark [487]: *Evolutionary Algorithms in Molecular Design*
- Chen [458]: *Evolutionary Computation in Economics and Finance*
- Ghosh and Jain [465]: *Evolutionary Computation in Data Mining*
- Miettinen, Mäkelä, Neittaanmäki, and Periaux [456]: *Evolutionary Algorithms in Engineering and Computer Science*
- Fogel [606]: *Evolutionary Computation: Principles and Practice for Signal Processing*
- Ashlock [607]: *Evolutionary Computation for Modeling and Optimization*
- Watanabe and Hashem [475]: *Evolutionary Computations – New Algorithms and their Applications to Evolutionary Robots*
- Cagnoni, Lutton, and Olague [476]: *Genetic and Evolutionary Computation for Image Processing and Analysis*
- Kramer [608]: *Self-Adaptive Heuristics for Evolutionary Computation*
- Lobo, Lima, and Michalewicz [609]: *Parameter Setting in Evolutionary Algorithms*
- Spears [610]: *Evolutionary Algorithms – The Role of Mutation and Recombination*
- Eiben and Michalewicz [611]: *Evolutionary Computation*
- Jin [612]: *Knowledge Incorporation in Evolutionary Computation*
- Grosan, Abraham, and Ishibuchi [613]: *Hybrid Evolutionary Algorithms*
- Abraham, Jain, and Goldberg [614]: *Evolutionary Multiobjective Optimization*
- Kallel, Naudts, and Rogers [615]: *Theoretical Aspects of Evolutionary Computing*
- Ghosh and Tsutsui [616]: *Advances in Evolutionary Computing – Theory and Applications*
- Yang, Shan, and Bui [617]: *Success in Evolutionary Computation*

2.3 Fitness Assignment

2.3.1 Introduction

With concept of prevalence comparisons introduced in Section 1.2.2 on page 9 we define a partial order on the elements in the problem space $\mathbb{X}$. Many selection algorithms however require that a total order is imposed onto the individuals p of a population. Such a total order can be created by assigning a single real number $v(p.x)$ to each solution candidate $p.x$ – its *fitness*.

The fitness assigned to an individual may not just reflect its rank in the population, but can also incorporate density/niching information. This way, not only the quality of a solution candidate is considered, but also the overall diversity of the population. This can improve the chance of finding the global optima as well as the performance of the optimization algorithm significantly. If many individuals in the population occupy the same rank or do not dominate each other, for instance, such information will be very helpful.

The fitness $v(p.x)$ thus not only depends on the solution candidate $p.x$ itself, but on the whole population Pop of the evolutionary algorithm and on the archive Arc of optimal elements, if available. In practical realizations, the fitness values are often stored in a special member variable in the individual records. Therefore, $v(p.x)$ can be considered as a mapping that returns the value of such a variable which has previously been stored there by a fitness assignment process “assignFitness”.

Definition 2.5 (Fitness Assignment). A fitness assignment process “assignFitness” creates a function $v : \mathbb{X} \mapsto \mathbb{R}^+$ which relates a scalar fitness value to each solution candidate in the population Pop Equation 2.1 (and archive Arc , if an archive is available Equation 2.2).

$$v = \text{assignFitness}(Pop, \text{cmp}_F) \Rightarrow v(p.x) \in \mathbb{V} \subseteq \mathbb{R}^+ \forall p \in Pop \quad (2.1)$$

$$v = \text{assignFitness}(Pop, Arc, \text{cmp}_F) \Rightarrow v(p.x) \in \mathbb{V} \subseteq \mathbb{R}^+ \forall p \in Pop \cup Arc \quad (2.2)$$

In the context of this book, we generally minimize fitness values, i. e., the lower the fitness of a solution candidate the better. Therefore, many of the fitness assignment processes based on the prevalence relation will obey to Equation 2.3. This equation represents a general relation – sometimes it is useful to violate it for some individuals in the population, especially when crowding information is incorporated.

$$p_1.x \succ p_2.x \Rightarrow v(p_1.x) < v(p_2.x) \quad \forall p_1, p_2 \in Pop \cup Arc \quad (2.3)$$

2.3.2 Weighted Sum Fitness Assignment

The most primitive fitness assignment strategy would be assigning a weighted sum of the objective values. This approach is very static and comes with the same problems as weighted sum-based approach for defining what an optimum is introduced in Section 1.2.3 on page 11. It makes no use of the prevalence relation. For computing the weighted sum of the different objective values of a solution candidate, we reuse Equation 1.4 on page 11 from the weighted sum optimum definition. The weights have to be chosen in a way that ensures that $v(p.x) \in \mathbb{R}^+$ holds for all individuals p .

$$v(p.x) = \text{assignFitnessWeightedSum}(Pop) \Leftrightarrow \forall p \in Pop \Rightarrow v(p.x) = g(p.x) \quad (2.4)$$

2.3.3 Pareto Ranking

Another very simple method of fitness assignment would be to use fitness values that directly reflect the prevalence relation. Figure 2.4 and Table 2.1 illustrate the Pareto relations in a population of 15 individuals and their corresponding objective values f_1 and f_2 , both subject to minimization. For doing so, we can imagine two different approaches:

To each individual, we can assign a value inversely proportional to the number of other individuals it prevails, like $v(p_1.x) \equiv \frac{1}{|\{p_2 \in Pop : p_1.x \succ p_2.x\}| + 1}$. We have written such fitness values in the column “Ap. 1” of Table 2.1 for Pareto optimization, i. e., the special case where the Pareto dominance relation is used to define prevalence. Individuals that dominate many others will here receive a lower fitness value than those which are prevailed by many.




Fig. 2.4: An example scenario for Pareto ranking.

x	prevails	is prevailed by	Ap. 1	Ap. 2
1	{5, 6, 8, 9, 14, 15}	$\emptyset$	$1/7$	0
2	{6, 7, 8, 9, 10, 11, 13, 14, 15}	$\emptyset$	$1/10$	0
3	{12, 13, 14, 15}	$\emptyset$	$1/5$	0
4	$\emptyset$	$\emptyset$	1	0
5	{8, 15}	{1}	$1/3$	1
6	{8, 9, 14, 15}	{1, 2}	$1/5$	2
7	{9, 10, 11, 14, 15}	{2}	$1/6$	1
8	{15}	{1, 2, 5, 6}	$1/2$	4
9	{14, 15}	{1, 2, 6, 7}	$1/3$	4
10	{14, 15}	{2, 7}	$1/3$	2
11	{14, 15}	{2, 7}	$1/3$	2
12	{13, 14, 15}	{3}	$1/4$	1
13	{15}	{2, 3, 12}	$1/2$	3
14	{15}	{1, 2, 3, 6, 7, 9, 10, 11, 12}	$1/2$	9
15	$\emptyset$	{1, 2, 3, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14}	1	13

Table 2.1: The Pareto domination relation of the individuals illustrated in Figure 2.4.

When taking a look at these values, the disadvantage of this approach becomes clear: It promotes individuals that reside in crowded region of the problem space and underrates those in sparsely explored areas. By doing so, it achieves exactly the opposite of what we want. Instead of exploring the problem space and delivering a wide scan of the frontier of best possible solution candidates, it will focus all effort on one single trajectory. Will we only obtain a subset of the best solutions and it is even possible that this fitness assignment method leads to premature convergence to a local optimum. A good example for this are the four non-prevailed individuals $\{1, 2, 3, 4\}$ from the Pareto frontier. The best fitness is assigned to the element 2, followed by individual 1. Although individual 7 is dominated (by 1), its fitness is better than the fitness of the non-dominated element 3. The solution candidate 4 gets the worst possible fitness 1, since it prevails no other element. Its chances for reproduction are similarly low than those of individual 15 which is dominated by all other elements except 4. Hence, both solution candidates will most probably be not selected and vanish in the next generation. The loss of solution candidate 4 will greatly decrease the diversity and even increase the focus on the crowded area near 1 and 2.

A much better approach for fitness assignment directly based on prevalence has first been proposed by Goldberg [139]. Here, the idea is to assign the number of individuals it is prevailed by to each solution candidate [618, 88, 619, 620]. This way, the previously mentioned negative effects will not occur. As the column “Ap 2” in Table 2.1 shows, all four non-prevailed individuals now have the best possible fitness 0. Hence, the exploration pressure is applied to a much wider area of the Pareto frontier. This so-called *Pareto ranking* can be performed by first removing all non-prevailed individuals from the population and assigning the rank 0 to them. Then, the same is performed with the rest of the population, and the individuals only dominated by those on rank 0 (now non-dominated) will be removed and get the rank 1. This is repeated until all solution candidates have a proper fitness assigned to them. Algorithm 2.3 outlines another simple way to perform Pareto ranking. Since we follow the idea of the freer prevalence comparators instead of Pareto dominance relations, we will synonymously refer to this approach as *Prevalence ranking*.

Algorithm 2.3: $v \leftarrow \text{assignFitnessParetoRank}(Pop, \text{cmp}_F)$

Input: Pop : the population to assign fitness values to
Input: cmp_F : the prevalence comparator defining the prevalence relation
Data: i, j, cnt : the counter variables
Output: v : a fitness function reflecting the Prevalence ranking

```

1 begin
2   for  $i \leftarrow \text{len}(Pop) - 1$  down to 0 do
3      $cnt \leftarrow 0$ 
4      $p \leftarrow Pop[i]$ 
5     for  $j \leftarrow \text{len}(Pop) - 1$  down to 0 do
6       // Check whether  $\text{cmp}_F(Pop[j].x, p.x) < 0$ 
7       if  $(j \neq i) \wedge (Pop[j].x \succ p.x)$  then  $cnt \leftarrow cnt + 1$ 
8      $v(p.x) \leftarrow cnt$ 
9   return  $v$ 
10 end

```

As already mentioned, the fitness values of all non-prevailed elements in our example Figure 2.4 and Table 2.1 are equally 0. However, the region around 1 and 2 has probably already exhaustively been explored, whereas the surrounding of 4 is rather unknown. A better approach of fitness assignment should incorporate such information and put a bit more pressure into the direction of 4, in order to make the evolutionary algorithm explore this area more thoroughly.

2.3.4 Sharing Functions

Previously, we have mentioned that the drawback of Pareto ranking is that it does not incorporate any information about whether the solution candidates reside closely to each other or in regions of the problem space which are only sparsely covered by individuals. Sharing, as a method for including such diversity information into the fitness assignment process, was introduced by Holland [253] and later refined by Deb [621], Goldberg and Richardson [622], Deb and Goldberg [623]. [624, 85]

Definition 2.6 (Sharing Function). A sharing function $\text{Sh} : \mathbb{R}^+ \mapsto \mathbb{R}^+$ is a function used to relate two individuals p_1 and p_2 to a value that decreases with their distance¹⁴ $d = \text{dist}(p_1, p_2)$ in a way that it is 1 for $d = 0$ and 0 if the distance exceeds a specified constant σ .

¹⁴ The concept of distance and a set of different distance measures is defined in Section 28.1 on page 469.

$$\text{Sh}_\sigma(d = \text{dist}(p_1, p_2)) = \begin{cases} 1 & \text{if } d \leq 0 \\ \text{Sh}_\sigma(d) \in [0, 1] & \text{if } 0 < d < \sigma \\ 0 & \text{otherwise} \end{cases} \quad (2.5)$$

Sharing functions can be employed in many different ways and are used by a variety of fitness assignment processes [622, 625]. Typically, the simple triangular function Sh_tri [626] or one of its either convex (Sh_cvexp) or concave (Sh_ccavp) pendants with the power $p \in \mathbb{R}^+, p > 0$ are applied. Besides using different powers of the distance- σ -ratio, another approach is the exponential sharing method Sh_exp .

$$\text{Sh\_tri}_\sigma(\sigma) d = \begin{cases} 1 - \frac{d}{\sigma} & \text{if } 0 \leq d < \sigma \\ 0 & \text{otherwise} \end{cases} \quad (2.6)$$

$$\text{Sh\_cvex}_{\sigma,p}(d) = \begin{cases} \left(1 - \frac{d}{\sigma}\right)^p & \text{if } 0 \leq d < \sigma \\ 0 & \text{otherwise} \end{cases} \quad (2.7)$$

$$\text{Sh\_ccav}_{\sigma,p}(d) = \begin{cases} 1 - \left(\frac{d}{\sigma}\right)^p & \text{if } 0 \leq d < \sigma \\ 0 & \text{otherwise} \end{cases} \quad (2.8)$$

$$\text{Sh\_exp}_{\sigma,p}(d) = \begin{cases} 1 & \text{if } d \leq 0 \\ 0 & \text{if } d \geq \sigma \\ \frac{e^{-\frac{pd}{\sigma}} - e^{-p}}{1 - e^{-p}} & \text{otherwise} \end{cases} \quad (2.9)$$

For sharing, we may use the distance of the individuals in the search space $\mathbb{G}$ as well as their distance in the problem space $\mathbb{X}$ or the objective space $\mathbb{Y}$. Often, the Euclidian distance measure $\text{dist}_{\text{eucl}}(F(p_1.x), F(p_2.x))$ in the objective space is used in sharing functions, but of course, any other distance measure could be applied. If the solution candidates are real vectors in the $\mathbb{R}^n$, we could as well measure the Euclidean distance of the phenotypes of the individuals directly, i. e., compute $\text{dist}_{\text{eucl}}(p_1.x, p_2.x)$. In genetic algorithms, where the search space is the set of all bit strings $\mathbb{G} = \mathbb{B}^n$ of the length n , another suitable approach would be to use the Hamming distance¹⁵ $\text{dist}_{\text{Ham}}(p_1.g, p_2.g)$ of the genotypes. The work of Deb [621], however, indicates that phenotypical sharing will often be superior to genotypical sharing.

Definition 2.7 (Niche Count). The niche count $m(p, P)$ [628, 624] of an individual p is the sum its sharing values with all other individual in a list P .

$$\forall p \in P \Rightarrow m(p, P) = \sum_{i=0}^{\text{len}(P)-1} \text{Sh}_\sigma(\text{dist}(p, P[i])) \quad (2.10)$$

The niche count m is always greater than zero, since $p \in P$ and, hence, $\text{Sh}_\sigma(\text{dist}(p, p)) = 1$ is computed and added up at least once.

The original sharing approach was developed for single-objective optimization where only one objective function f was subject to maximization. In this case, it was simply divided by the niche count, punishing solutions in crowded regions [624]. The goal of sharing was to distribute the population over a number of different peaks in the fitness landscape, with each peak receiving a fraction of the population proportional to its height [626]. The results of dividing the fitness by the niche counts strongly depends on the height differences of the peaks and thus, on the complexity class¹⁶ of f . On $f_1 \in \mathbf{O}(n)$, for instance, the influence of m is much bigger than on a $f_2 \in \mathbf{O}(e^n)$.

By multiplying the niche count m to predetermined fitness values v' , we can use this approach for fitness minimization in conjunction with a variety of other different fitness assignment processes, but also inherit its shortcomings:

¹⁵ See Definition 28.6 on page 469 for more information on the Hamming distance.

¹⁶ See Section 29.1.3 on page 482 for a detailed introduction into complexity and the $\mathbf{O}$ -notation.

$$v(p.x) = v'(p.x) * m(p, Pop), \quad v' \equiv \text{assignFitness}(Pop, \text{cmp}_F) \quad (2.11)$$

Sharing was traditionally combined with fitness proportionate, i. e., roulette wheel selection¹⁷. Oei et al. [85] have shown that if the sharing function is computed using the parental individuals of the “old” population and then naïvely combined with the more sophisticated tournament selection¹⁸, the resulting behavior of the evolutionary algorithm may be chaotic. They suggested to use the partially filled “new” population to circumvent this problem. The layout of evolutionary algorithms, as defined in this book, bases the fitness computation on the whole set of “new” individuals and assumes that their objective values have already been completely determined. In other words, such issues simply do not exist in multi-objective evolutionary algorithms as introduced here and the chaotic behavior does occur.

For computing the niche count m , $\mathcal{O}(n^2)$ comparisons are needed. According to Goldberg et al. [629], sampling the population can be sufficient to approximate m in order to avoid this quadratic complexity.

2.3.5 Variety Preserving Ranking

Using sharing and the niche counts plainly leads to more or less unpredictable effects. Of course, it promotes solutions located in sparsely populated niches but how much their fitness will be improved is rather unclear. Using distance measures which are not normalized can lead to strange effects too. Imagine two objective functions f_1 and f_2 . If the values of f_1 span from 0 to 1 for the individuals in the population whereas those of f_2 range from 0 to 10000, the components of f_1 will most often be negligible in the Euclidian distance of two individuals in the objective space $\mathbb{Y}$. Another problem is that the effect of simple sharing on the pressure into the direction of the Pareto frontier is not obvious either or depends on the sharing approach applied. Some methods simply add a niche count to the Pareto rank, which may cause non-dominated individuals having worse fitness than any others in the population. Other approaches scale the niche count into the interval $[0, 1)$ before adding it, which not only ensures that non-dominated individuals have the best fitness but also leave the relation between individuals at different ranks intact, which does not further variety very much.

Variety Preserving Ranking is a fitness assignment approach based on Pareto ranking using prevalence comparators and sharing. We have developed it in order to mitigate all these previously mentioned side effects and balance the evolutionary pressure between optimizing the objective functions and maximizing the variety inside the population. In the following, we will describe the process of Variety Preserving Ranking-based fitness assignment which is defined in Algorithm 2.4.

Before the process can begin, it is required that all individuals with infinite objective values must be removed from the population Pop . If such a solution candidate is optimal, i. e., if it has negative infinitely large objectives in a minimization process, for instance, it should receive fitness zero, since fitness is subject to minimization. If the individual is infeasible, on the other hand, its fitness should be set to $\text{len}(Pop) + \sqrt{\text{len}(Pop)} + 1$, which is larger by one than every other fitness values that may be assigned by Algorithm 2.4.

In lines 2 to 9, we create a list *rank*s which we use to efficiently compute the Pareto rank of every solution candidate in the population. By the way, the word *prevalence rank* would be more precise in this case, since we use prevalence comparisons as introduced in Section 1.2.7. Therefore, Variety Preserving Ranking is not limited to Pareto optimization but may also incorporate External Decision Makers (Section 1.2.6) or the method of inequalities (Section 1.2.5). The highest rank encountered in the population is stored in the variable *maxRank*. This value may be zero if the population contains only non-prevalent elements. The lowest rank will always be zero since the prevalence comparators cmp_F define

¹⁷ Roulette wheel selection is discussed in Section 2.4.3 on page 95.

¹⁸ You can find an outline of tournament selection in Section 2.4.4 on page 100.

Algorithm 2.4: $v \leftarrow \text{assignFitnessVariancePreserving}(Pop, \text{cmp}_F)$

Input: Pop : the population
Input: cmp_F : the comparator function
Input: $[implicit] F$: the set of objective functions
Data: ...: sorry, no space here, we'll discuss this in the text
Output: v : the fitness function

```

1 begin
  /* If needed: Remove all elements with infinite objective values from Pop
    and assign fitness 0 or  $\text{len}(Pop) + \sqrt{\text{len}(Pop)} + 1$  to them. Then compute the
    prevalence ranks. */
2  ranks  $\leftarrow \text{createList}(\text{len}(Pop), 0)$ 
3  maxRank  $\leftarrow 0$ 
4  for  $i \leftarrow \text{len}(Pop) - 1$  down to 0 do
5    for  $j \leftarrow i - 1$  down to 0 do
6       $k \leftarrow \text{cmp}_F(Pop[i].x, Pop[j].x)$ 
7      if  $k < 0$  then ranks[j]  $\leftarrow$  ranks[j] + 1
8      else if  $k > 0$  then ranks[i]  $\leftarrow$  ranks[i] + 1
9      if ranks[i] > maxRank then maxRank  $\leftarrow$  ranks[i]
  // determine the ranges of the objectives
10 mins  $\leftarrow \text{createList}(|F|, +\infty)$ 
11 maxs  $\leftarrow \text{createList}(|F|, -\infty)$ 
12 foreach  $p \in Pop$  do
13   for  $i \leftarrow |F|$  down to 1 do
14     if  $f_i(p.x) < \text{mins}[i-1]$  then mins[i-1]  $\leftarrow f_i(p.x)$ 
15     if  $f_i(p.x) > \text{maxs}[i-1]$  then maxs[i-1]  $\leftarrow f_i(p.x)$ 
16 rangeScales  $\leftarrow \text{createList}(|F|, 1)$ 
17 for  $i \leftarrow |F| - 1$  down to 0 do
18   if maxs[i] > mins[i] then rangeScales[i]  $\leftarrow 1 / (\text{maxs}[i] - \text{mins}[i])$ 
  // Base a sharing value on the scaled Euclidean distance of all elements
19 shares  $\leftarrow \text{createList}(\text{len}(Pop), 0)$ 
20 minShare  $\leftarrow +\infty$ 
21 maxShare  $\leftarrow -\infty$ 
22 for  $i \leftarrow \text{len}(Pop) - 1$  down to 0 do
23   curShare  $\leftarrow \text{shares}[i]$ 
24   for  $j \leftarrow i - 1$  down to 0 do
25     dist  $\leftarrow 0$ 
26     for  $k \leftarrow |F|$  down to 1 do
27       dist  $\leftarrow \text{dist} + [(f_k(Pop[i].x) - f_k(Pop[j].x)) * \text{rangeScales}[k-1]]^2$ 
28      $s \leftarrow \text{Sh\_exp}_{\sqrt{|F|}, 16}(\sqrt{\text{dist}})$ 
29     curShare  $\leftarrow \text{curShare} + s$ 
30     shares[j]  $\leftarrow \text{shares}[j] + s$ 
31   shares[i]  $\leftarrow \text{curShare}$ 
32   if curShare < minShare then minShare  $\leftarrow$  curShare
33   if curShare > maxShare then maxShare  $\leftarrow$  curShare
  // Finally, compute the fitness values
  scale  $\leftarrow \begin{cases} 1 / (\text{maxShare} - \text{minShare}) & \text{if } \text{maxShare} > \text{minShare} \\ 1 & \text{otherwise} \end{cases}$ 
34 for  $i \leftarrow \text{len}(Pop) - 1$  down to 0 do
35   if ranks[i] > 0 then
36      $v(Pop[i].x) \leftarrow \text{ranks}[i] + \sqrt{\text{maxRank}} * \text{scale} * (\text{shares}[i] - \text{minShare})$ 
37   else  $v(Pop[i].x) \leftarrow \text{scale} * (\text{shares}[i] - \text{minShare})$ 
38
39 end

```

order relations which are non-circular by definition.¹⁹ We will use *maxRank* to determine the maximum penalty for solutions in an overly crowded region of the search space later on.

From line 10 to 18, we determine the maximum and the minimum values that each objective function takes on when applied to the individuals in the population. These values are used to store the inverse of their ranges in the array *rangeScales*, which we will use to scale all distances in each dimension (objective) of the individuals into the interval $[0, 1]$. There are $|F|$ objective functions in F and, hence, the maximum Euclidian distance between two solution candidates in the (scaled) objective space becomes $\sqrt{|F|}$. It occurs if all the distances in the single dimensions are 1.

The most complicated part of the Variety Preserving Ranking algorithm is between line 19 and 33. Here we computed the scaled distance from every individual to each other solution candidate in the objective space and use this distance to aggregate share values (in the array *shares*). Therefore, again two nested loops are needed (lines 22 and 24). The distance components of two individuals $Pop[i].x$ and $Pop[j].x$ are scaled summarized in a variable *dist* in line 27. The Euclidian distance between them is $\sqrt{dist}$ which we use to determine a sharing value in 28. We therefore have decided for exponential sharing with power 16 and $\sigma = \sqrt{|F|}$, as introduced in Equation 2.9 on page 86. For every individual, we sum up all the shares (see line 30). While doing so, we also determine the minimum and maximum such total share in the variables *minShare* and *maxShare* in lines 32 and 33.

We will use these variables to scale all sharing values again into the interval $[0, 1]$ (line 34), so the individual in the most crowded region always has a total share of 1 and the most remote individual always has a share of 0. So basically, we now know two things about the individuals in *Pop*:

1. their Pareto rank, stored in the array *ranks*, giving information about their relative quality according to the objective values and
2. their sharing values, held in *shares*, denoting how densely crowded the area around them is.

With this information, we determine the final fitness values of an individual p as follows: If p is non-prevalled, i. e., its rank is zero, its fitness is its scaled total share (line 38). Otherwise, we multiply the square root of the maximum rank, $\sqrt{maxRank}$, with the scaled share and add it to its rank (line 37). By doing so, we preserve the supremacy of non-prevalled individuals in the population but allow them to compete with each other based on the crowdedness of their location in the objective space. All other solution candidates may degenerate in rank, but at most by the square root of the worst rank.

Example

Let us now apply Variety Preserving Ranking to the examples for Pareto ranking from Section 2.3.3. In Table 2.2, we again list all the solution candidates from Figure 2.4 on page 84, this time with their objective values obtained with f_1 and f_2 corresponding to their coordinates in the diagram. In the third column, you can find the Pareto rank of the individuals as it has been listed in Table 2.1 on page 84. The columns *share/u* and *share/s* correspond to the total sharing sums of the individuals, unscaled and scaled into $[0, 1]$.

But let us start from the beginning. As already mentioned, we know the Pareto ranks of the solution candidates from Table 2.1, so the next step is to determine the ranges of values the objective functions take on for the example population. These can easily be found out from Figure 2.4. f_1 spans from 1 to 10, which leads to $rangeScale[0] = 1/9$. $rangeScale[1] = 1/8$ since the maximum of f_2 is 9 and its minimum is 1. With this, we now can compute the (dimensional scaled) distances amongst the solution candidates in the objective space, the values of $\sqrt{dist}$ in algorithm Algorithm 2.4 as well as the corresponding

¹⁹ In all order relations imposed on finite sets always result in at least one “smallest” element. See Section 26.7.2 on page 415 for more information




Fig. 2.5: The sharing potential in the Variety Preserving Ranking example

x	f_1	f_2	rank	share/u	share/s	$v(x)$
1	1	7	0	0.71	0.779	0.779
2	2	4	0	0.239	0.246	0.246
3	6	2	0	0.201	0.202	0.202
4	10	1	0	0.022	0	0
5	1	8	1	0.622	0.679	3.446
6	2	7	2	0.906	1	5.606
7	3	5	1	0.531	0.576	3.077
8	2	9	4	0.314	0.33	5.191
9	3	7	4	0.719	0.789	6.845
10	4	6	2	0.592	0.645	4.325
11	5	5	2	0.363	0.386	3.39
12	7	3	1	0.346	0.366	2.321
13	8	4	3	0.217	0.221	3.797
14	7	7	9	0.094	0.081	9.292
15	9	9	13	0.025	0.004	13.01

Table 2.2: An example for Variety Preserving Ranking based on Figure 2.4.

values of the sharing function $\text{Sh\_exp}_{\sqrt{|F|},16}(\sqrt{\text{dist}})$. We have noted these in Table 2.3, using the upper triangle of the table for the distances and the lower triangle for the shares.

Upper triangle: distances. Lower triangle: corresponding share values.															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1		0.391	0.836	1.25	0.125	0.111	0.334	0.274	0.222	0.356	0.51	0.833	0.863	0.667	0.923
2	0.012		0.51	0.965	0.512	0.375	0.167	0.625	0.391	0.334	0.356	0.569	0.667	0.67	0.998
3	7.7E-5	0.003		0.462	0.933	0.767	0.502	0.981	0.708	0.547	0.391	0.167	0.334	0.635	0.936
4	6.1E-7	1.8E-5	0.005		1.329	1.163	0.925	1.338	1.08	0.914	0.747	0.417	0.436	0.821	1.006
5	0.243	0.003	2.6E-5	1.8E-7		0.167	0.436	0.167	0.255	0.417	0.582	0.914	0.925	0.678	0.898
6	0.284	0.014	1.7E-4	1.8E-6	0.151		0.274	0.25	0.111	0.255	0.417	0.747	0.765	0.556	0.817
7	0.023	0.151	0.003	2.9E-5	0.007	0.045		0.512	0.25	0.167	0.222	0.51	0.569	0.51	0.833
8	0.045	0.001	1.5E-5	1.5E-7	0.151	0.059	0.003		0.274	0.436	0.601	0.933	0.914	0.609	0.778
9	0.081	0.012	3.3E-4	4.8E-6	0.056	0.284	0.059	0.045		0.167	0.334	0.669	0.67	0.444	0.712
10	0.018	0.023	0.002	3.2E-5	0.009	0.056	0.151	0.007	0.151		0.167	0.502	0.51	0.356	0.67
11	0.003	0.018	0.012	2.1E-4	0.001	0.009	0.081	0.001	0.023	0.151		0.334	0.356	0.334	0.669
12	8E-5	0.002	0.151	0.009	3.2E-5	2.1E-4	0.003	2.6E-5	0.001	0.003	0.023		0.167	0.5	0.782
13	5.7E-5	0.001	0.023	0.007	2.9E-5	1.7E-4	0.002	3.2E-5	0.001	0.003	0.018	0.151		0.391	0.635
14	0.001	0.001	0.001	9.3E-5	4.6E-4	0.002	0.003	0.001	0.007	0.018	0.023	0.003	0.012		0.334
15	2.9E-5	1.2E-5	2.5E-5	1.1E-5	3.9E-5	9.7E-5	8E-5	1.5E-4	3.1E-4	0.001	0.001	1.4E-4	0.001	0.023	

Table 2.3: The distance and sharing matrix of the example from Table 2.2.

You also can imagine sharing as a scalar field, as illustrated in Figure 2.5. In this case, imagine each individual in the population as an electron that will build an electrical field around it resulting in a potential. If two electrons come close, repulsing forces occur, which is pretty much the same what we want to do with Variety Preserving Ranking. Unlike the electrical field, the power of the sharing potential falls exponentially, resulting in relatively steep spikes in Figure 2.5 which gives proximity and density a heavier influence. Electrons in atoms on planets are limited in their movement by other influences like gravity or nuclear forces, which are often stronger than the electromagnetic force. In Variety Preserving Ranking, the prevalence rank plays this role – as you can see in Table 2.2, its influence on the fitness is often dominant.

But back to the example. Summing up the single share potentials for each individual, we obtain the fifth column of Table 2.3, the unscaled share values. Their minimum is around 0.022 and the maximum is 0.94. Therefore, we must subtract 0.022 from each of these values and multiply the result with 1.131. By doing so, we build the column *shares/s*. Finally, we can compute the fitness values according to lines 38 and 37 in Algorithm 2.4.

The last column of Table 2.2 lists these results. As you can see, all non-prevalent individuals have retained a fitness value less than one, lower than those of any other solution candidate in the population. However, amongst these best individuals, solution candidate 4 is strongly preferred, since it is located in a very remote location of the objective space. Individual 1 is the least interesting non-dominated one, because it has the densest neighborhood in Figure 2.4. In this neighborhood, the individuals 5 and 6 with the Pareto ranks 1 and 2 are located. They are strongly penalized by the sharing process and receive the fitness values $v(5) = 3.446$ and $v(6) = 5.606$. In other words, individual 5 becomes less interesting than solution candidate 7 which has a worse Pareto rank. 6 now is even worse than individual 8 which would have a fitness better by two if strict Pareto ranking was applied.

Based on these fitness values, algorithms like Tournament selection (see Section 2.4.2) or fitness proportionate approaches (discussed in Section 2.4.3) will pick elements in a way that preserves the pressure into the direction of the Pareto frontier but also leads to a balanced and sustainable variety in the population.

2.3.6 Tournament Fitness Assignment

In tournament fitness assignment, which is a generalization of the q -level binary tournament selection introduced by Weicker [255], the fitness of each individual is computed by letting it compete q times against r other individuals (with $r = 1$ as default) and counting the number of competitions it loses. For a better understanding of the tournament metaphor see Section 2.4.4 on page 100, where the tournament selection scheme is discussed. Anyway, the number of losses will approximate its Pareto rank, but are a bit more randomized than that. If we would count the number of tournaments won instead of the losses, we would encounter the same problems than in the first idea of Pareto ranking.

Algorithm 2.5: $v \leftarrow \text{assignFitnessTournament}_{q,r}(Pop, \text{cmp}_F)$

Input: q : the number of tournaments per individuals
Input: r : the number of other contestants per tournament, normally 1
Input: Pop : the population to assign fitness values to
Input: cmp_F : the comparator function providing the prevalence relation
Data: i, j, k, z : counter variables
Data: b : a Boolean variable being **true** as long as a tournament isn't lost
Data: p : the individual currently examined
Output: v : the fitness function

```

1 begin
2   for  $i \leftarrow \text{len}(Pop) - 1$  down to 0 do
3      $z \leftarrow q$ 
4      $p \leftarrow Pop[i]$ 
5     for  $j \leftarrow q$  down to 1 do
6        $b \leftarrow \text{true}$ 
7        $k \leftarrow r$ 
8       while  $(k > 0) \wedge b$  do
9          $b \leftarrow Pop[\text{random}_u(0, \text{len}(Pop))] \succ p.x$ 
10         $k \leftarrow k - 1$ 
11      if  $b$  then  $z \leftarrow z - 1$ 
12     $v(p.x) \leftarrow z$ 
13  return  $v$ 
14 end
```

**TODO add remaining fitness
assignment methods**

2.4 Selection

2.4.1 Introduction

Definition 2.8 (Selection). In evolutionary algorithms, the selection²⁰ operation $Mate = \text{select}(Pop, v, ms)$ chooses ms individuals according to their fitness values v from the population Pop and places them into the mating pool $Mate$ [409, 57, 630, 55].

²⁰ http://en.wikipedia.org/wiki/Selection_%28genetic_algorithm%29 [accessed 2007-07-03]

$$\begin{aligned}
Mate = \text{select}(Pop, v, ms) \Rightarrow & \forall p \in Mate \Rightarrow p \in Pop \\
& \forall p \in Pop \Rightarrow p \in \mathbb{G} \times \mathbb{X} \\
& v(p.x) \in \mathbb{R}^+ \quad \forall p \in Pop \\
& (\text{len}(Mate) \geq \min\{\text{len}(Pop), ms\}) \wedge (\text{len}(Mate) \leq ms)
\end{aligned} \tag{2.12}$$

On the mating pool, the reproduction operations discussed in Section 2.5 on page 108 will subsequently be applied. Selection may behave in a deterministic or in a randomized manner, depending on the algorithm chosen and its application-dependant implementation. Furthermore, elitist evolutionary algorithms may incorporate an archive *Arc* in the selection process, as sketched in Algorithm 2.2.

Generally, there are two classes of selection algorithms: such *with replacement* (annotated with a subscript $_r$) and such *without replacement* (annotated with a subscript $_w$, see Equation 2.13) [631]. In a selection algorithm *without replacement*, each individual from the population *Pop* is taken into consideration for reproduction at most once and therefore also will occur in the mating pool *Mate* one time at most. The mating pool returned by algorithms *with replacement* can contain the same individual multiple times. Like in nature, one individual may thus have multiple offspring. Normally, selection algorithms are used in a variant with replacement. One of the reasons therefore is the number of elements to be placed into the mating pool (corresponding to the parameter *ms*). If $\text{len}(Pop) < ms$, the mating pool returned by a method without replacement contains less than *ms* individuals since it can at most consist of the whole population.

$$Mate = \text{select}_w(Pop, v, ms) \Rightarrow \text{countItemOccurences}(p, Mate) = 1 \quad \forall p \in Mate \tag{2.13}$$

The selection algorithms have major impact on the performance of evolutionary algorithms. Their behavior has thus been subject to several detailed studies, conducted by, for instance, Goldberg and Deb [632], Blickle and Thiele [630], and Zhong et al. [633], just to name a few.

Usually, fitness assignment processes are carried out before selection and the selection algorithms base their decisions solely on the fitness *v* of the individuals. It is possible to rely on the prevalence relation, i. e., to write $\text{select}(Pop, \text{cmp}_F, ms)$ instead of $\text{select}(Pop, v, ms)$, thus saving the costs of the fitness assignment process. However, this will lead to the same problems that occurred in the first approach of prevalence-proportional fitness assignment (see Section 2.3.3 on page 83) and we will therefore not discuss such techniques in this book.

Many selection algorithms only work with scalar fitness and thus need to rely on a fitness assignment process in multi-objective optimization. Selection algorithms can be chained – the resulting mating pool of the first selection may then be used as input for the next one, maybe even with a secondary fitness assignment process in between. In some applications, an environmental selection that reduces the number of individuals is performed first and then a mating selection follows which extracts the individuals which should be used for reproduction.

Visualization

In the following sections, we will discuss multiple selection algorithms. In order to ease understanding them, we will visualize the expected number of times $S(p)$ that an individual *p* will reach the mating pool *Mate* for some of the algorithms.

$$S(p) = E[\text{countItemOccurences}(p, Mate)] \tag{2.14}$$

Therefore, we will use the special case where we have a population *Pop* of $\text{len}(Pop) = 1000$ individuals, $p_0..p_{999}$ and also a target mating pool size $ms = 1000$. Each individual p_i has the fitness value $v(p_i.x)$, and fitness is subject to minimization. For this fitness, we consider two cases:

1. As sketched in Fig. 2.6.1, the individual p_i has fitness i , i. e., $v_1(p_0.x) = 0, v_1(p_1.x) = 1, \dots, v_1(p_{999}.x) = 999$.
2. Individual p_i has fitness $(i + 1)^3$, i. e., $v_2(p_0.x) = 1, v_2(p_1.x) = 3, \dots, v_2(p_{999}.x) = 1\,000\,000\,000$, as illustrated in Fig. 2.6.2.



Fig. 2.6.1: Case 1: $v(p_i.x) = i$ Fig. 2.6.2: Case 2: $v(p_i.x) = (i + 1)^3$

Fig. 2.6: The two example fitness cases.

2.4.2 Truncation Selection

Truncation selection²¹, also called deterministic selection or threshold selection, returns the $k < ms$ best elements from the list Pop . These elements are copied as often as needed until the mating pool size ms reached. For k , normally values like $\lceil \text{len}(Pop)/2 \rceil$ or $\lceil \text{len}(Pop)/3 \rceil$ are used. Algorithm 2.6 realizes this scheme by first sorting the population in ascending order according to the fitness v . Then, it iterates from 0 to ms and inserts only the elements with indices from 0 to $k - 1$ into the mating pool.

Algorithm 2.6: $Mate \leftarrow \text{truncationSelect}_k(Pop, v, ms)$

Input: Pop : the list of individuals to select from

Input: v : the fitness values

Input: ms : the number of individuals to be placed into the mating pool $Mate$

Input: k : cut-off value

Data: i : counter variables

Output: $Mate$: the winners of the tournaments which now form the mating pool

```

1 begin
2    $Mate \leftarrow ()$ 
3    $k \leftarrow \min\{k, \text{len}(Pop)\}$ 
4    $Pop \leftarrow \text{sortList}_a(Pop, v)$ 
5   for  $i \leftarrow 0$  up to  $ms - 1$  do
6      $Mate \leftarrow \text{addListItem}(Mate, Pop[i \bmod k])$ 
7   return  $Mate$ 
8 end
```

Truncation selection is usually used in Evolution Strategies with $(\mu + \lambda)$ and (μ, λ) strategies. In general evolutionary algorithms, it should be combined with a fitness assignment

²¹ http://en.wikipedia.org/wiki/Truncation_selection [accessed 2007-07-03]

process that incorporates diversity information in order to prevent premature convergence. Recently, Lässig et al. [634] have proved that truncation selection is the optimal selection strategy for crossover, provided that the right value of k is used. In practical applications, this value is normally not known.




Fig. 2.7: The number of expected offspring in truncation selection.

In Figure 2.7, we sketch the expected number of offspring for the individuals from our examples specified in Section 2.4.1. In this selection scheme, the diagram will look exactly the same regardless whether we use fitness configuration 1 or 2, since it is solely based on the order of individuals and not on the numerical relation of their fitness. If we set $k = ms = \text{len}(Pop)$, each individual will have one offspring in average. If $k = \frac{1}{2}ms$, the top-50% individuals will have two offspring and the others none. For $k = \frac{1}{10}ms$, only the best 100 from the 1000 solution candidates will reach the mating pool but reproduce 10 times in average.

2.4.3 Fitness Proportionate Selection

Fitness proportionate selection²² has already been applied in the original genetic algorithms as introduced by Holland [253] and therefore is one of the oldest selection schemes. In fitness proportionate selection, the probability $P(p_1)$ of an individual $p_1 \in Pop$ to enter the mating pool is proportional to its fitness $v(p.x)$ (subject to maximization) compared to the sum of the fitness of all individuals. This relation in its original form is defined in Equation 2.15 below.

$$P(p_1) = \frac{v(p_1.x)}{\sum_{p_2 \in Pop} v(p_2.x)} \quad (2.15)$$

There exists a variety of approaches which realize such probability distributions [632], like stochastic remainder selection (Brindle [635], Booker [636]) and stochastic universal

²² http://en.wikipedia.org/wiki/Fitness_proportionate_selection [accessed 2008-03-19]

selection (Baker [637], Greffentette and Baker [417]). The most commonly known method is the Monte Carlo *roulette wheel selection* by De Jong [254], where we imagine the individuals of a population to be placed on a roulette²³ wheel as sketched in Fig. 2.8.1. The size of the area on the wheel standing for a solution candidate is proportional to its fitness. The wheel is spun, and the individual where it stops is placed into the mating pool *Mate*. This procedure is repeated until *ms* individuals have been selected.




Fig. 2.8.1: Example for fitness maximization.




Fig. 2.8.2: Example for normalized fitness minimization.

Fig. 2.8: Examples for the idea of roulette wheel selection.

In the context of this book, fitness is subject to minimization. Here, higher fitness values $v(p.x)$ indicate unfit solution candidates $p.x$ whereas lower fitness denotes high utility. Furthermore, the fitness values are normalized into a range of $[0, sum]$, because otherwise, fitness proportionate selection will handle the set of fitness values $\{0, 1, 2\}$ in a different way than $\{10, 11, 12\}$. Equation 2.19 defines the framework for such a (normalized) fitness proportionate selection “*rouletteWheelSelect*”. It is illustrated exemplarily in Fig. 2.8.2 and realized in Algorithm 2.7 as a variant with and in Algorithm 2.8 without replacement. Amongst others, Whitley [446] points out that even fitness normalization as performed here cannot overcome the drawbacks of fitness proportional selection methods.

$$\min V = \min \{v(p.x) \mid \forall p \in Pop\} \quad (2.16)$$

$$\max V = \max \{v(p.x) \mid \forall p \in Pop\} \quad (2.17)$$

$$\text{norm } V(p.x) = \frac{\max V - v(p.x)}{\max V - \min V} \quad (2.18)$$

$$P(p_1) = \frac{\text{norm } V(p_1.x)}{\sum_{\forall p_2 \in Pop} \text{norm } V(p_2.x)} \quad (2.19)$$

But what are the drawbacks of fitness proportionate selection methods? Let us therefore visualize the expected results of roulette wheel selection applied to the special cases stated in Section 2.4.1. Figure 2.9 illustrates the number of expected occurrences $S(p_i)$ of an individual p_i if roulette wheel selection was applied. Since $ms = 1000$, we draw one thousand times a single individual from the population *Pop*. Each single choice is based on the proportion of the individual fitness in the total fitness of all individuals, as defined in Equation 2.15 and Equation 2.19. Thus, in scenario 1 with the fitness sum $\frac{999 \cdot 998}{2} = 498501$, the relation

²³ <http://en.wikipedia.org/wiki/Roulette> [accessed 2008-03-20]

Algorithm 2.7: $Mate \leftarrow \text{rouletteWheelSelect}_r(Pop, v, ms)$

Input: Pop : the list of individuals to select from
Input: v : the fitness values
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Data: i : a counter variable
Data: a : a temporary store for a numerical value
Data: A : the array of fitness values
Data: min, max, sum : the minimum, maximum, and sum of the fitness values
Output: $Mate$: the mating pool

```

1 begin
2    $A \leftarrow \text{createList}(\text{len}(Pop), 0)$ 
3    $min \leftarrow \infty$ 
4    $max \leftarrow -\infty$ 
5   for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
6      $a \leftarrow v(Pop[i].x)$ 
7      $A[i] \leftarrow a$ 
8     if  $a < min$  then  $min \leftarrow a$ 
9     if  $a > max$  then  $max \leftarrow a$ 
10  if  $max = min$  then
11     $max \leftarrow max + 1$ 
12     $min \leftarrow min - 1$ 
13   $sum \leftarrow 0$ 
14  for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
15     $sum \leftarrow \frac{max - A[i]}{max - min}$ 
16     $A[i] \leftarrow sum$ 
17  for  $i \leftarrow 0$  up to  $ms - 1$  do
18     $a \leftarrow \text{searchItem}_{as}(\text{random}_u(0, sum), A)$ 
19    if  $a < 0$  then  $a \leftarrow -a - 1$ 
20     $Mate \leftarrow \text{addListItem}(Mate, Pop[a])$ 
21  return  $Mate$ 
22 end
  
```

$S(p_i) = ms * \frac{i}{498501}$ holds for fitness maximization and $S(p_i) = ms \frac{999-i}{498501}$ for minimization. As result (sketched in Fig. 2.9.1), the fittest individuals produce (on average) two offspring, whereas the worst solution candidates will always vanish in this example. For the 2nd scenario with $v_2(p_i.x) = (i + 1)^3$, the total fitness sum is approximately $2.51 \cdot 10^{11}$ and $S(p_i) = ms \frac{(i+1)^3}{2.52 \cdot 10^{11}}$ holds for maximization. The resulting expected values depicted in Fig. 2.9.2 are significantly different from those in Fig. 2.9.1. The meaning of this is that the design of the objective functions (or the fitness assignment process) has a much stronger influence on the convergence behavior of the evolutionary algorithm. This selection method only works well if the fitness of an individual is indeed something like a proportional measure for the probability that it will produce better offspring.

Thus, roulette wheel selection has a bad performance compared to other schemes like tournament selection [632, 638] or ranking selection [632, 630]. It is mainly included here for the sake of completeness and because it is easy to understand and suitable for educational purposes.

Algorithm 2.8: $Mate \leftarrow \text{rouletteWheelSelect}_w(Pop, v, ms)$

Input: Pop : the list of individuals to select from
Input: v : the fitness values
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Data: i : a counter variable
Data: a, b : temporary stores for numerical values
Data: A : the array of fitness values
Data: min, max, sum : the minimum, maximum, and sum of the fitness values
Output: $Mate$: the mating pool

```

1 begin
2    $A \leftarrow \text{createList}(\text{len}(Pop), 0)$ 
3    $min \leftarrow \infty$ 
4    $max \leftarrow -\infty$ 
5   for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
6      $a \leftarrow v(Pop[i].x)$ 
7      $A[i] \leftarrow a$ 
8     if  $a < min$  then  $min \leftarrow a$ 
9     if  $a > max$  then  $max \leftarrow a$ 
10  if  $max = min$  then
11     $max \leftarrow max + 1$ 
12     $min \leftarrow min - 1$ 
13   $sum \leftarrow 0$ 
14  for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
15     $sum \leftarrow \frac{max - A[i]}{max - min}$ 
16     $A[i] \leftarrow sum$ 
17  for  $i \leftarrow 0$  up to  $\min\{ms, \text{len}(Pop)\} - 1$  do
18     $a \leftarrow \text{searchItem}_{as}(\text{random}_u(0, sum), A)$ 
19    if  $a < 0$  then  $a \leftarrow -a - 1$ 
20    if  $a = 0$  then  $b \leftarrow 0$ 
21    else  $b \leftarrow A[a-1]$ 
22     $b \leftarrow A[a] - b$ 
23    for  $j \leftarrow a + 1$  up to  $\text{len}(A) - 1$  do
24       $A[j] \leftarrow A[j] - b$ 
25     $sum \leftarrow sum - b$ 
26     $Mate \leftarrow \text{addListItem}(Mate, Pop[a])$ 
27     $Pop \leftarrow \text{deleteListItem}(Pop, a)$ 
28     $A \leftarrow \text{deleteListItem}(A, a)$ 
29  return  $Mate$ 
30 end

```



Fig. 2.9.1: $v_1(p_i.x) = i$ Fig. 2.9.2: $v_1(p_i.x) = (i + 1)^3$

Fig. 2.9: The number of expected offspring in roulette wheel selection.

2.4.4 Tournament Selection

Tournament selection²⁴, proposed by Wetzel [639] and studied by Brindle [635], is one of the most popular and effective selection schemes. Its features are well-known and have been analyzed by a variety of researchers such as Blickle and Thiele [638; 630], Miller and Goldberg [220], Lee et al. [640], Sastry and Goldberg [631], and Oei et al. [85]. In tournament selection, k elements are picked from the population *Pop* and compared with each other in a tournament. The winner of this competition will then enter mating pool *Mate*. Although being a simple selection strategy, it is very powerful and therefore used in many practical applications [641, 642, 643, 644].

As example, consider a tournament selection (with replacement) with a tournament size of two [645]. For each single tournament, the contestants are chosen randomly according to a uniform distribution and the winners will be allowed to enter the mating pool. If we assume that the mating pool will contain about as same as many individuals as the population, each individual will, on average, participate in two tournaments. The best solution candidate of the population will win all the contests it takes part in and thus, again on average, contributes approximately two copies to the mating pool. The median individual of the population is better than 50% of its challengers but will also lose against 50%. Therefore, it will enter the mating pool roughly one time on average. The worst individual in the population will lose all its challenges to other solution candidates and can only score even if competing against itself, which will happen with probability $(1/ms)^2$. It will not be able to reproduce in the average case because $ms * (1/ms)^2 = 1/ms < 1 \forall ms > 1$.

For visualization purposes, let us go back to our examples from Section 2.4.1 with a population of 1000 individuals $p_0..p_{999}$ and $ms = 1000$. Again, we assume that each individual has a unique fitness value of $v_1(p_i.x) = i$ or $v_2(p_i.x) = (i + 1)^3$, respectively. If we apply tournament selection with replacement in this special scenario, the expected number of occurrences $S(p_i)$ of an individual p_i in the mating pool can be computed according to Blickle and Thiele [630] as

$$S(p_i) = ms * \left(\left(\frac{1000 - i}{1000} \right)^k - \left(\frac{1000 - i - 1}{1000} \right)^k \right) \quad (2.20)$$

The absolute values of the fitness play no role. The only thing that matters is whether or not the fitness of one individual is higher as the fitness of another one, not fitness difference itself. The expected numbers of offspring for the two example cases 1 and 2 from Section 2.4.1 are the same. Tournament selection thus gets rid of the problems of fitness proportionate methods. Figure 2.10 depicts these numbers for different tournament sizes $k = \{1, 2, 3, 4, 5, 10\}$. If $k = 1$, tournament selection degenerates to randomly picking individuals and each solution candidate will occur one time in the mating pool on average. With rising k , the selection pressure increases: individuals with good fitness values create more and more offspring whereas the chance of worse solution candidates to reproduce decreases.

Tournament selection with replacement (TSR) is presented in Algorithm 2.9. Tournament selection without replacement (TSoR) [640, 646] can be defined in two forms. In the first variant specified as Algorithm 2.10, a solution candidate cannot compete against itself. This method is defined in. In Algorithm 2.11, on the other hand, an individual may enter the mating pool at most once.

The algorithms specified here should more precisely be entitled as *deterministic* tournament selection algorithms since the winner of the k contestants that take part in each tournament enters the mating pool. In the non-deterministic variant this is not necessarily the case. There, a probability p is defined. The best individual in the tournament is selected with probability p , the second best with probability $p(1 - p)$, the third best with probability $p(1 - p)^2$ and so on. The i^{th} best individual in a tournament enters the mating pool with

²⁴ http://en.wikipedia.org/wiki/Tournament_selection [accessed 2007-07-03]




Fig. 2.10: The number of expected offspring in tournament selection.

Algorithm 2.9: $Mate \leftarrow \text{tournamentSelect}_{r,k}(Pop, v, ms)$

Input: Pop : the list of individuals to select from**Input:** v : the fitness values**Input:** ms : the number of individuals to be placed into the mating pool $Mate$ **Input:** [implicit] k : the tournament size**Data:** a : the index of the tournament winner**Data:** i, j : counter variables**Output:** $Mate$: the winners of the tournaments which now form the mating pool

```

1 begin
2    $Mate \leftarrow ()$ 
3    $Pop \leftarrow \text{sortList}_a(Pop, v)$ 
4   for  $i \leftarrow 0$  up to  $ms - 1$  do
5      $a \leftarrow \lfloor \text{random}_u(0, \text{len}(Pop)) \rfloor$ 
6     for  $j \leftarrow 1$  up to  $k - 1$  do
7        $a \leftarrow \min \{a, \lfloor \text{random}_u(0, \text{len}(Pop)) \rfloor\}$ 
8      $Mate \leftarrow \text{addListItem}(Mate, Pop[a])$ 
9   return  $Mate$ 
10 end

```

Algorithm 2.10: $Mate \leftarrow \text{tournamentSelect}_{w1,k}(Pop, v, ms)$

Input: Pop : the list of individuals to select from
Input: v : the fitness values
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Input: k : the tournament size
Data: a : the index of the tournament winner
Data: i, j : counter variables
Output: $Mate$: the winners of the tournaments which now form the mating pool

```

1 begin
2    $Mate \leftarrow ()$ 
3    $Pop \leftarrow \text{sortList}_a(Pop, v)$ 
4   for  $i \leftarrow 0$  up to  $\min\{\text{len}(Pop), ms\} - 1$  do
5      $a \leftarrow \lfloor \text{random}_u(0, \text{len}(Pop)) \rfloor$ 
6     for  $j \leftarrow 1$  up to  $\min\{\text{len}(Pop), k\} - 1$  do
7        $a \leftarrow \min\{a, \lfloor \text{random}_u(0, \text{len}(Pop)) \rfloor\}$ 
8      $Mate \leftarrow \text{addListItem}(Mate, Pop[a])$ 
9      $Pop \leftarrow \text{deleteListItem}(Pop, a)$ 
10  return  $Mate$ 
11 end
  
```

Algorithm 2.11: $Mate \leftarrow \text{tournamentSelect}_{w2,k}(Pop, v, ms)$

Input: Pop : the list of individuals to select from
Input: v : the fitness values
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Input: k : the tournament size
Data: A : the list of contestants per tournament
Data: a : the tournament winner
Data: i, j : counter variables
Output: $Mate$: the winners of the tournaments which now form the mating pool

```

1 begin
2    $Mate \leftarrow ()$ 
3    $Pop \leftarrow \text{sortList}_a(Pop, v)$ 
4   for  $i \leftarrow 0$  up to  $ms - 1$  do
5      $A \leftarrow ()$ 
6     for  $j \leftarrow 1$  up to  $\min\{k, \text{len}(Pop)\}$  do
7       repeat
8          $a \leftarrow \lfloor \text{random}_u(0, \text{len}(Pop)) \rfloor$ 
9         until  $\text{searchItem}_u(a, A) < 0$ 
10       $A \leftarrow \text{addListItem}(A, a)$ 
11     $a \leftarrow \min A$ 
12     $Mate \leftarrow \text{addListItem}(Mate, Pop[a])$ 
13  return  $Mate$ 
14 end
  
```

probability $p(1-p)^i$. Algorithm 2.12 realizes this behavior for a tournament selection with replacement. Notice that it becomes equivalent to Algorithm 2.9 on page 101 if p is set to 1. Besides the algorithms discussed here, a set of additional tournament-based selection methods has been introduced by Lee et al. [640].

Algorithm 2.12: $Mate \leftarrow \text{tournamentSelect}_{r,k}^p(Pop, v, ms)$

Input: Pop : the list of individuals to select from
Input: v : the fitness values
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Input: [implicit] p : the selection probability, $p \in [0, 1]$
Input: [implicit] k : the tournament size
Data: A : the set of tournament contestants
Data: i, j : counter variables
Output: $Mate$: the winners of the tournaments which now form the mating pool

```

1 begin
2    $Mate \leftarrow ()$ 
3    $Pop \leftarrow \text{sortList}_a(Pop, v)$ 
4   for  $i \leftarrow 0$  up to  $ms - 1$  do
5      $A \leftarrow ()$ 
6     for  $j \leftarrow 0$  up to  $k - 1$  do
7        $A \leftarrow \text{addListItem}(A, \lfloor \text{random}_u(0, \text{len}(Pop)) \rfloor)$ 
8      $A \leftarrow \text{sortList}_a(A, \text{cmp}(a_1, a_2) \equiv (a_1 - a_2))$ 
9     for  $j \leftarrow 0$  up to  $\text{len}(A) - 1$  do
10      if  $(\text{random}_u() \leq p) \vee (j \geq \text{len}(A) - 1)$  then
11         $Mate \leftarrow \text{addListItem}(Mate, Pop[A[j]])$ 
12         $j \leftarrow \infty$ 
13   return  $Mate$ 
14 end

```

2.4.5 Ordered Selection

Ordered selection is another approach for circumventing the problems of fitness proportionate selection methods. Here, the probability of an individual to be selected is proportional to (a power of) its position (rank) in the sorted list of all individuals in the population. The implicit parameter $k \in \mathbb{R}^+$ of the ordered selection algorithm determines the selection pressure. It equals to the number of expected offspring of the best individual and is thus much similar to the parameter k of tournament selection. The bigger k gets, the higher is the probability that individuals which are non-prevalled i.e., have good objective values will be selected.

Algorithm 2.13 demonstrates how ordered selection with replacement works and the variant without replacement is described in Algorithm 2.14. Basically, it first converts the parameter k to a power q to which the uniformly drawn random numbers are raised that are used for indexing the sorted individual list. This can be achieved with Equation 2.21.

$$q = \frac{1}{1 - \frac{\log k}{\log ms}} \quad (2.21)$$

Figure 2.11 illustrates the expected offspring in the application of ordered selection with $k \in \{1, 2, 3, 4, 5\}$. Like tournament selection, a value of $k = 1$ leads degenerates the evolutionary algorithm to a parallel random walk. Another close similarity to tournament selection

Algorithm 2.13: $Mate \leftarrow \text{orderedSelect}_r^p(Pop, v, ms)$

Input: Pop : the list of individuals to select from
Input: v : the fitness values
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Input: k : the parameter of the ordering selection
Data: q : the power value to be used for ordering
Data: i : a counter variable
Output: $Mate$: the mating pool

```

1 begin
2    $q \leftarrow \frac{1}{1 - \frac{\log k}{\log ms}}$ 
3    $Mate \leftarrow ()$ 
4    $Pop \leftarrow \text{sortList}_a(Pop, v)$ 
5   for  $i \leftarrow 0$  up to  $ms - 1$  do
6      $Mate \leftarrow \text{addListItem}(Mate, Pop[\lfloor \text{random}_u()^p * \text{len}(Pop) \rfloor])$ 
7   return  $Mate$ 
8 end

```

Algorithm 2.14: $Mate \leftarrow \text{orderedSelect}_w^p(Pop, v, ms)$

Input: Pop : the list of individuals to select from
Input: v : the fitness values
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Input: k : the parameter of the ordering selection
Data: q : the power value to be used for ordering
Data: i, j : counter variables
Output: $Mate$: the mating pool

```

1 begin
2    $q \leftarrow \frac{1}{1 - \frac{\log k}{\log ms}}$ 
3    $Mate \leftarrow ()$ 
4    $Pop \leftarrow \text{sortList}_a(Pop, v)$ 
5   for  $i \leftarrow 0$  up to  $\min\{ms, \text{len}(Pop)\} - 1$  do
6      $j \leftarrow \lfloor \text{random}_u()^p * \text{len}(Pop) \rfloor$ 
7      $Mate \leftarrow \text{addListItem}(Mate, Pop[j])$ 
8      $Pop \leftarrow \text{deleteListItem}(Pop, j)$ 
9   return  $Mate$ 
10 end

```

occurs when comparing the exact formulas computing the expected offspring for our examples:

$$S(p_i) = ms * \left(\left(\frac{i+1}{1000} \right)^q - \left(\frac{i}{1000} \right)^q \right) \quad (2.22)$$

Equation 2.22 looks pretty much like Equation 2.20. The differences between the two selection methods become obvious when comparing the diagrams Figure 2.11 and Figure 2.10 which both are independent of the actual fitness values. Tournament selection creates many copies of the better fraction of the population and almost none of the others. Ordered selection focuses on an even smaller group of the fittest individuals but also even the worst solution candidates still have a survival probability not too far from one. In other words, while tournament selection reproduces a larger group of good individuals and kills most of the others, ordered selection assigns very high fertility to very few individuals but preserves also the less fitter ones.




Fig. 2.11: The number of expected offspring in ordered selection.

2.4.6 Ranking Selection

Ranking selection, introduced by Baker [647] and more thoroughly discussed by Whitley [446], Bickel and Thiele [630; 648], and Goldberg and Deb [632] is another approach for circumventing the problems of fitness proportionate selection methods. In ranking selection [647, 446, 417], the probability of an individual to be selected is proportional to its position (rank) in the sorted list of all individuals in the population. Using the rank smoothes out larger differences of the objective values and emphasizes small ones. Generally, we can the conventional ranking selection method as the application of a fitness assignment process setting the rank as fitness (which can be achieved with Pareto ranking) and a subsequent fitness proportional selection.

2.4.7 VEGA Selection

The *Vector Evaluated Genetic Algorithm* by Schaffer [649; 650] applies a special selection algorithm which does not incorporate any preceding fitness assignment process but works on the objective values directly. For each of the objective functions $f_i \in F$, it selects a subset of the mating pool *Mate* of the size $ms/|F|$. Therefore it applies fitness proportionate selection which is based on f_i instead of a fitness assignment “assignFitness”. The mating pool is then a mixture of these sub-selections. Richardson et al. [651] show in [652] that this selection scheme is approximately the same as if computing a weighted sum of the fitness values. As pointed out by Fonseca and Fleming [653], in the general case, this selection method will sample non-prevalled solution candidates at different frequencies. Schaffer also anticipated that the population of his GA may split into different species, each particularly strong in one objective, if the Pareto frontier is concave.

Algorithm 2.15: $Mate \leftarrow \text{vegaSelectionPopFms}$

Input: Pop : the list of individuals to select from
Input: F : the objective functions
Input: ms : the number of individuals to be placed into the mating pool $Mate$
Data: i : a counter variable
Data: j : the size of the current subset of the mating pool
Data: A : a temporary mating pool
Output: $Mate$: the individuals selected

```

1 begin
2    $Mate \leftarrow ()$ 
3   for  $i \leftarrow 1$  up to  $|F|$  do
4      $j \leftarrow \frac{ms}{|F|}$ 
5     if  $i = 1$  then  $j \leftarrow j + ms \bmod |F|$ 
6      $A \leftarrow \text{rouletteWheelSelect}_r(Pop, v \equiv f_i, j)$ 
7      $Mate \leftarrow \text{appendList}(Mate, A)$ 
8   return  $Mate$ 
9 end
  
```

2.4.8 Simple Convergence Prevention

In our experiments (especially in Genetic Programming and problems with discrete objective functions) we often use a very simple mechanism to prevent premature convergence (see Section 1.4.1) which we outline in Algorithm 2.16. It is neither a fitness nor a selection algorithm, but I think it fits best into this section. The idea is simple: the more similar individuals we have in the population, the more likely are we converged. We do not know whether we have converged to a global optimum or to a local one. If we got stuck at a local optimum, we should maybe limit the fraction of the population which resides at this spot. In case we have found the global optimum, this approach does not hurt, because, in the end, one single point on this optimum suffices. In order to measure similarity, we use the most general measure that the evolutionary algorithm has: the objective functions. All individuals are compared with each other. If two have exactly the same objective values, one of them is thrown away with probability $cp \in [0, 1]$ and also does not take part in further comparisons. This way, we weed out similar individuals and make room in the population and mating pool for a wider diversity of solution candidates. For $cp = 0$, this prevention mechanism is turned off, for $cp = 1$, all remaining individuals will have different objective values.

Although this approach is almost hilariously simple, the results of our experiments were often significantly better with this convergence prevention method turned on than without it. Additionally, in none of our experiments, the outcomes were influenced negatively by this filter, which makes it even more robust than other methods for convergence prevention like sharing or variety preserving.

If an individual p occurs n times in the population or if there are n individuals with exactly the same objective values, Algorithm 2.16 cuts down the expected number of their occurrences $S(p)$ to

$$S(p) = \sum_{i=1}^n (1 - cp)^{i-1} = \sum_{i=0}^{n-1} (1 - cp)^i = \frac{(1 - cp)^n - 1}{-cp} = \frac{1 - (1 - cp)^n}{cp} \quad (2.23)$$

In Figure 2.12, we sketch the expected number of remaining instances of the individual p after this pruning process, if it occurred n times in the population before Algorithm 2.16 was applied.

From Equation 2.23 follows that even a population of infinite size which has fully converged to one single value will probably not contain more than $\frac{1}{cp}$ copies of this individual

Algorithm 2.16: $Pop' \leftarrow \text{convergencePrevention}(Pop, cp)$

Input: Pop : the list of individuals to apply convergence prevention to
Input: cp : the convergence prevention probability, $cp \in [0, 1]$
Input: $[implicit] F$: the list of objective functions
Data: i, j : counter variables
Data: p : the individual checked in this generation
Output: Pop' : the pruned population

```

1 begin
2    $Pop' \leftarrow ()$ 
3   for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
4      $p \leftarrow Pop[i]$ 
5     for  $j \leftarrow \text{len}(Pop') - 1$  down to 0 do
6       if  $f(p.x) = f(Pop'[j].x) \ \forall f \in F$  then
7         if  $\text{random}_u() < cp$  then
8            $Pop' \leftarrow \text{deleteListItem}(Pop', j)$ 
9        $Pop' \leftarrow \text{addListItem}(Pop', p)$ 
10  return  $Pop'$ 
11 end
  
```




Fig. 2.12: The expected numbers of occurrences for different values of n and cp .

after the simple convergence prevention has been applied. This threshold is also visible in Figure 2.12.

$$\lim_{n \rightarrow \infty} S(p) = \lim_{n \rightarrow \infty} \frac{1 - (1 - cp)^n}{cp} = \frac{1 - 0}{cp} = \frac{1}{cp} \quad (2.24)$$

TODO add remaining selection algorithms

2.5 Reproduction

An optimization algorithm uses the information gathered up to step t for creating the solution candidates to be evaluated in step $t + 1$. There exist different methods to do so. In evolutionary algorithms, the aggregated information corresponds to the population Pop and the set of best individuals Arc if such an archive is maintained. The search operations $searchOp \in Op$ in used in the evolutionary algorithm family are called *reproduction* operation, inspired by the biological procreation mechanisms²⁵ of mother nature [416]. There are four basic operations:

- *Creation* has no direct natural paragon; it simply creates a new genotype without any ancestors or heritage. Hence, it roughly can be compared with the occurrence of the first living cells from out a soup of certain chemicals²⁶.
- *Duplication* resembles the cell division²⁷, resulting in two individuals similar to one parent.
- *Mutation* in evolutionary algorithms corresponds to small, random variations in the genotype of an individual, exactly like its natural counterpart²⁸.
- Like in sexual reproduction, *recombination*²⁹ combines two parental genotypes to a new genotype including traits from both elders.

In the following, we will discuss these operations in detail and provide general definitions for them.

Definition 2.9 (Creation). The creation operation “create” is used to produce a new genotype $g \in \mathbb{G}$ with a random configuration.

$$g = \text{create}() \Rightarrow g \in \mathbb{G} \quad (2.25)$$

When an evolutionary algorithm starts, no information about the search space has been gathered yet. Hence, we cannot use existing solution candidates to derive new ones and search operations with an arity higher than zero cannot be applied. Creation is thus used to fill the initial population $Pop(t = 0)$.

Definition 2.10 (Duplication). The duplication operation $\text{duplicate} : \mathbb{G} \mapsto \mathbb{G}$ is used to create an exact copy of an existing genotype $g \in \mathbb{G}$.

$$g = \text{duplicate}(g) \quad \forall g \in \mathbb{G} \quad (2.26)$$

Duplication is just a placeholder for copying an element of the search space, i. e., it is what occurs when neither mutation nor recombination are applied. It is useful to increase the share of a given type of individual in a population.

Definition 2.11 (Mutation). The mutation operation $\text{mutate} : \mathbb{G} \mapsto \mathbb{G}$ is used to create a new genotype $g_n \in \mathbb{G}$ by modifying an existing one. The way this modification is performed is application-dependent. It may happen in a randomized or in a deterministic fashion.

$$g_n = \text{mutate}(g) : g \in \mathbb{G} \Rightarrow g_n \in \mathbb{G} \quad (2.27)$$

Definition 2.12 (Recombination). The recombination (or crossover³⁰) operation $\text{recombine} : \mathbb{G} \times \mathbb{G} \mapsto \mathbb{G}$ is used to create a new genotype $g_n \in \mathbb{G}$ by combining the features of two existing ones. Depending on the application, this modification may happen in a randomized or in a deterministic fashion.

$$g_n = \text{recombine}(g_a, g_b) : g_a, g_b \in \mathbb{G} \Rightarrow g_n \in \mathbb{G} \quad (2.28)$$

²⁵ <http://en.wikipedia.org/wiki/Reproduction> [accessed 2007-07-03]

²⁶ <http://en.wikipedia.org/wiki/Abiogenesis> [accessed 2008-03-17]

²⁷ http://en.wikipedia.org/wiki/Cell_division [accessed 2008-03-17]

²⁸ <http://en.wikipedia.org/wiki/Mutation> [accessed 2007-07-03]

²⁹ http://en.wikipedia.org/wiki/Sexual_reproduction [accessed 2008-03-17]

³⁰ <http://en.wikipedia.org/wiki/Recombination> [accessed 2007-07-03]

Notice that the term *recombination* is more general than *crossover* since it stands for arbitrary search operations that combines the traits of two individuals. Crossover, however, is only used if the elements search space $\mathbb{G}$ are linear representations. Then, it stands for exchanging parts of these so-called strings.

Now we can define the set Op_{EA} of search operations most commonly applied in evolutionary algorithms as

$$Op_{EA} = \{\text{create, duplicate, mutate, recombine}\} \quad (2.29)$$

All of them can be combined arbitrarily. It is, for instance, not unusual to mutate the results of a recombination operation, i. e., to perform $\text{mutate}(\text{recombine}(g_1, g_2))$.

The four operators are altogether used to reproduce whole populations of individuals.

Definition 2.13 (*reproducePop*). The population reproduction operation $Pop = \text{reproducePop}(Mate)$ is used to create a new population Pop by applying the reproduction operations to the mating pool $Mate$.

$$\begin{aligned} Pop = \text{reproducePop}(Mate) \Rightarrow & \forall p \in Mate \Rightarrow p \in P, \forall p \in Pop \Rightarrow p \in P, \text{len}(Pop) = \text{len}(Mate) \\ & \forall p \in Pop \Rightarrow p.g = \text{create}() \vee \\ & \quad p.g = \text{duplicate}(p_{old}.g) : p_{old} \in Mate \vee \\ & \quad p.g = \text{mutate}(p_{old}.g) : p_{old} \in Mate \vee \\ & \quad p.g = \text{recombine}(p_{old1}.g, p_{old2}.g) : \\ & \quad \quad p_{old1}, p_{old2} \in Mate \end{aligned} \quad (2.30)$$

For creating an initial population of the size s , we furthermore define the function $\text{createPop}(s)$ in Algorithm 2.17.

Algorithm 2.17: $Pop \leftarrow \text{createPop}(s)$

Input: s : the number of individuals in the new population

Input: [implicit] create : the creation operator

Data: i : a counter variable

Output: Pop : the new population of randomly created individuals ($\text{len}(Pop) = s$)

```

1 begin
2    $Pop \leftarrow ()$ 
3   for  $i \leftarrow 0$  up to  $s - 1$  do
4      $Pop \leftarrow \text{addListItem}(Pop, \text{create}())$ 
5   return  $Pop$ 
6 end

```

2.5.1 NCGA Reproduction

The Neighborhood Cultivation Genetic Algorithm by Watanabe et al. [654] discussed in ?? uses a special reproduction method. Recombination is performed only on neighboring individuals, which leads to child genotypes close to their parents. This so-called neighborhood cultivation shifts the recombination-operator more into the direction exploitation, i. e., NCGA uses crossover for investigating the close surrounding of known solution candidates. The idea is that parents that do not differ much from each other are more likely to be compatible in order to produce functional offspring than parents that have nothing in common. Neighborhood cultivation is achieved in Algorithm 2.18 by sorting the mating pool along one *focused* objective. Then, the elements situated directly besides each other are recombined. The focus on the objective rotates in a way that in a three-objective optimization the first

objective is focused at the beginning, then the second, then the third and after that again the first. The algorithm shown here receives the additional parameter foc which denotes the focused objective. Both, recombination and mutation are performed with an implicitly defined probability (r and m , respectively).

Algorithm 2.18: $Pop \leftarrow \text{ncgaReproducePop}_{foc}(Mate)$

Input: $Mate$: the mating pool

Input: foc : the objective currently focused

Input: [implicit] recombine , mutate : the recombination and mutation routines

Input: [implicit] r, m : the probabilities of recombination and mutation

Data: i : a counter variable

Output: Pop : the new population with $\text{len}(Pop) = \text{len}(Mate)$

```

1 begin
2    $Pop \leftarrow \text{sortList}_a(Mate, f_{foc})$ 
3   for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
4     if  $(\text{random}_u() \leq r) \wedge (i < \text{len}(Pop) - 1)$  then  $Pop[i] \leftarrow \text{recombine}(Pop[i], Pop[i+1])$ 
5     if  $\text{random}_u() \leq m$  then  $Pop[i] \leftarrow \text{mutate}(Pop[i])$ 
6   return  $Pop$ 
7 end

```

2.6 Algorithms

Besides the basic evolutionary algorithms introduced in Section 2.1.1 on page 71, there exists a variety of other, more sophisticated approaches. Many of them deal especially with multi-objective optimization which imposes new challenges on fitness assignment and selection. In this section we discuss the most prominent of these evolutionary algorithms.

2.6.1 VEGA

The very first multi-objective genetic algorithm is the *Vector Evaluated Genetic Algorithm* (VEGA) created by Schaffer [649; 650] in the mid-1980s. The main difference between VEGA and the basic form of evolutionary algorithms is the modified selection algorithm which you can find discussed in Section 2.4.7 on page 105. This selection algorithm solely relies on the objective functions F and does not use any preceding fitness assignment process nor can it incorporate a prevalence comparison scheme cmp_F . However, it has severe weaknesses also discussed in Section 2.4.7 and thus cannot be considered as an efficient approach to multi-objective optimization.

Algorithm 2.19: $X^* \leftarrow \text{vega}(F, s)$

Input: F : the objective functions**Input:** ps : the population size**Data:** t : the generation counter**Data:** Pop : the population**Data:** $Mate$: the mating pool**Data:** v : the fitness function resulting from the fitness assigning process**Output:** X^* : the set of the best elements found

```

1 begin
2    $t \leftarrow 0$ 
3    $Pop \leftarrow \text{createPop}(ps)$ 
4   while  $\text{terminationCriterion}()$  do
5      $Mate \leftarrow \text{vegaSelect}(Pop, F, ps)$ 
6      $t \leftarrow t + 1$ 
7      $Pop \leftarrow \text{reproducePop}(Mate)$ 
8   return  $\text{extractPhenotypes}(\text{extractOptimalSet}(Pop))$ 
9 end

```

TODO add remaining EAs

Genetic Algorithms

3.1 Introduction

Genetic algorithms¹ (GAs) are a subclass of evolutionary algorithms where the elements of the search space $\mathbb{G}$ are usually binary strings ($\mathbb{G} = \mathbb{B}^*$) or arrays of other elementary types. As sketched in Figure 3.1, the genotypes are used in the reproduction operations whereas the values of the objective functions $f \in F$ are computed on basis of the phenotypes in the problem space $\mathbb{X}$ which are obtained via the genotype-phenotype mapping “gpm” [1, 645].

The roots of genetic algorithms go back to the mid-1950s, where biologists like Barricelli [655; 656; 657; 658] and the computer scientist Fraser [659] began to apply computer-aided simulations in order to gain more insight into genetic processes and the natural evolution and selection. Bremermann [660] and Bledsoe [661; 662; 663; 664] used evolutionary approaches based on binary string genomes for solving inequalities, for function optimization, and for determining the weights in neural networks (Bledsoe and Browning [665]) in the early 1960s. At the end of that decade, important research on such search spaces was contributed by Bagley [666], Cavicchio, Jr. [667], and Frantz [668] – all conducted by Holland at the University of Michigan. As a result of Holland’s work, genetic algorithms as a new approach for problem solving finally became widely recognized and popular [669, 253, 670, 139]. Today, there are many applications in science, economy, and research and development [527] that can be tackled with genetic algorithms. Therefore, various forms of genetic algorithms [671] have been developed to. Some genetic algorithms² like the human-based genetic algorithms³ (HBGA), for instance, even require human beings for evaluating or selecting the solution candidates [672, 673, 674, 675, 676]

It should further be mentioned that, because of the close relation to biology and since genetic algorithms were originally applied to single-objective optimization, the objective functions f here are often referred to as *fitness functions*. This is a historically grown misnaming which should not be mixed up with the fitness assignment processes discussed in Section 2.3 on page 82 and the fitness values v used in the context of this book.

¹ http://en.wikipedia.org/wiki/Genetic_algorithm [accessed 2007-07-03]

² http://en.wikipedia.org/wiki/Interactive_genetic_algorithm [accessed 2007-07-03]

³ <http://en.wikipedia.org/wiki/HBGA> [accessed 2007-07-03]




Fig. 3.1: The basic cycle of genetic algorithms.

3.2 General Information

3.2.1 Areas Of Application

Some example areas of application of genetic algorithms are:

Application	References
Scheduling	[677, 678, 679, 680, 681, 682, 683]
Chemistry, Chemical Engineering	[684, 685, 686, 687, 688, 689, 690, 691]
Medicine	[692, 693, 694, 695]
Data Mining and Data Analysis	[696, 697, 698, 699, 464]
Geometry and Physics	[700, 701, 702, 703, 704]
Economics and Finance	[705]
Networking and Communication	[706, 707, 708, 709, 710, 711, 712, 713, 714, 715, 716, 717, 484, 718, 719, 720, 721, 722, 723, 724]

Electrical Engineering and Circuit Design	[725, 726, 727]
Image Processing	[728]
Combinatorial Optimization	[729, 730]

3.2.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on genetic algorithms are:

EUROGEN: Evolutionary Methods for Design Optimization and Control with Applications to Industrial Problems
see Section 2.2.2 on page 78

FOGA: Foundations of Genetic Algorithms

<http://www.sigevo.org/> [accessed 2007-09-01]

History: 2007: Mexico City, México, see [731]

2005: Aizu-Wakamatsu City, Japan, see [732]

2002: Torremolinos, Spain, see [733]

2000: Charlottesville, VA, USA, see [734]

1998: Madison, WI, USA, see [735]

1996: San Diego, CA, USA, see [736]

1994: Estes Park, Colorado, USA, see [737]

1992: Vail, Colorado, USA, see [738]

1990: Bloomington Campus, Indiana, USA, see [188]

FWGA: Finnish Workshop on Genetic Algorithms and Their Applications

NWGA: Nordic Workshop on Genetic Algorithms

History: 1997: Helsinki, Finland, see [739]

1996: Vaasa, Finland, see [740]

1995: Vaasa, Finland, see [741]

1994: Vaasa, Finland, see [742]

1992: Espoo, Finland, see [743]

GALESIA: International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications

now part of *CEC*, see Section 2.2.2 on page 77

History: 1997: Glasgow, UK, see [744]

1995: Scheffeld, UK, see [745]

GECCO: Genetic and Evolutionary Computation Conference

see Section 2.2.2 on page 79

ICGA: International Conference on Genetic Algorithms

Now part of *GECCO*, see Section 2.2.2 on page 79

History: 1997: East Lansing, Michigan, USA, see [746]

1995: Pittsburgh, PA, USA, see [747]

1993: Urbana-Champaign, IL, USA, see [748]

1991: San Diego, CA, USA, see [749]

1989: Fairfax, Virginia, USA, see [652]

1987: Cambridge, MA, USA, see [750]

1985: Pittsburgh, PA, USA, see [751]

ICANNGA: International Conference on Adaptive and Natural Computing Algorithms

see Section 2.2.2 on page 79

Mendel: International Conference on Soft Computing

see Section 1.6.2 on page 63

3.2.3 Online Resources

Some general, online available resources on genetic algorithms are:

<http://www.obitko.com/tutorials/genetic-algorithms/> [accessed 2008-05-17]

Last update: 1998

Description: A very thorough introduction to genetic algorithms by Marek Obitko

<http://www.aaii.org/AITopics/html/genalg.html> [accessed 2008-05-17]

Last update: up-to-date

Description: The genetic algorithms and Genetic Programming pages of the AAAI

<http://www.illigal.uiuc.edu/web/> [accessed 2008-05-17]

Last update: up-to-date

Description: The Illinois Genetic Algorithms Laboratory (IlliGAL)

<http://www.cs.cmu.edu/Groups/AI/html/faqs/ai/genetic/top.html> [accessed 2008-05-17]

Last update: 1997-08-10

Description: The Genetic Algorithms FAQ.

<http://www.rennard.org/alife/english/gavintrgb.html> [accessed 2008-05-17]

Last update: 2007-07-10

Description: An introduction to genetic algorithms by Jean-Philippe Rennard.

<http://www.optiwater.com/GAsearch/> [accessed 2008-06-08]

Last update: 2003-11-15

Description: GA-Search – The Genetic Algorithms Search Engine

3.2.4 Books

Some books about (or including significant information about) genetic algorithms are:

Goldberg [139]: *Genetic Algorithms in Search, Optimization and Machine Learning*

Mitchell [55]: *An Introduction to Genetic Algorithms*

Davis [752]: *Handbook of Genetic Algorithms*

Haupt and Haupt [753]: *Practical Genetic Algorithms*

Gen and Cheng [754]: *Genetic Algorithms and Engineering Design*

Chambers [755]: *Practical Handbook of Genetic Algorithms: Applications*

Chambers [756]: *Practical Handbook of Genetic Algorithms: New Frontiers*

Chambers [757]: *Practical Handbook of Genetic Algorithms: Complex Coding Systems*

Holland [253]: *Adaptation in Natural and Artificial Systems*

Gen and Chen [758]: *Genetic Algorithms (Engineering Design and Automation)*

Cant'u-Paz [759]: *Efficient and Accurate Parallel Genetic Algorithms*

Heistermann [760]: *Genetische Algorithmen. Theorie und Praxis evolutionärer Optimierung*

Schöneburg, Heinzmann, and Feddersen [761]: *Genetische Algorithmen und Evolutionsstrategien*

Gwiazda [762]: *Crossover for single-objective numerical optimization problems*

Schaefer and Telega [763]: *Foundations of Global Genetic Optimization*

Karr and Freeman [764]: *Industrial Applications of Genetic Algorithms*

Bäck [409]: *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*

Davis [765]: *Genetic Algorithms and Simulated Annealing*

3.3 Genomes in Genetic Algorithms

The search spaces $\mathbb{G}$ of genetic algorithms are usually referred to *genome*. As mentioned in Section 1.3.1 on page 23, its elements, the genotypes, encompass the whole hereditary information of an organism encoded in the DNA⁴. The DNA is a string of base pairs that encodes the phenotypical characteristics of the creature it belongs to. Like their natural prototypes, the genomes in genetic algorithms are strings, linear sequences of certain data types [139, 80, 55]. Because of the linear structure, these genotypes are also often called *chromosomes*. In genetic algorithms, we most often use chromosomes that are strings of one and the same data type, for example bits or real numbers.

Definition 3.1 (String Chromosome). A string chromosome can either be a fixed-length tuple (Equation 3.1) or a variable-length list (Equation 3.2).

In the first case, the loci i of the genes g_i are constant and, hence, the tuples may contain elements of different types $\mathbb{G}_i$.

$$\mathbb{G} = \{\forall (g[1], g[2], \dots, g[n]) : g[i] \in \mathbb{G}_i \forall i \in 1..n\} \quad (3.1)$$

This is not given in variable-length string genomes. Here, the positions of the genes may shift when the reproduction operations are applied. Thus, all elements of such genotypes must have the same type $\mathbb{G}_T$.

$$\mathbb{G} = \{\forall \text{lists } g : g[i] \in \mathbb{G}_T \forall 0 \leq i < \text{len}(g)\} \quad (3.2)$$

String chromosomes are often bit strings or vectors of real numbers. Genetic algorithms with such real vector genomes in their *natural representation* are called *real-encoded* [766]. In the case of bit string genomes, sometimes gray coding⁵ is used during the genotype-phenotype mapping. This by doing so, we try to preserve locality and ensure that small changes in the genotype will also lead to small changes in the phenotypes [767].

When reproducing sexually, the genotypes of the two parents genotypes will recombine. Additionally, small variations (mutations) will modify the chromosomes during this process. In asexual reproduction, mutations are the only changes that occur. The reproduction operations in genetic algorithms exactly copy the patterns of crossover and mutation.

After the genetic information has been reproduced, life begins with a single cell which divides⁶ time and again until a mature individual is formed⁷. The emergence of a phenotype $x \in \mathbb{X}$ from its genotypic representation $g \in \mathbb{G}$ is called *embryogenesis*, its counterparts in evolutionary search are the genotype-phenotype mapping and artificial embryogeny discussed in Section 3.6 on page 122.

Let us shortly recapitulate the structure of the elements g of the search space $\mathbb{G}$. A gene (see Definition 1.23 on page 24) is the basic informational unit in a genotype g . Depending on the genome, a gene can be a bit, a real number, or any other structure. In biology, a gene is a segment of nucleic acid that contains the information necessary to produce a functional RNA product in a controlled manner. An allele (see Definition 1.24) is a value of specific gene. A gene which is a bit, for example, can have the values $\{\text{true}, \text{false}\}$, whereas a gene representing a real number can take on alleles in $\mathbb{R}$, and so on. Thus, as in nature, an allele is a specific instance of a gene. The locus (see Definition 1.25) is the position where a specific gene can be found in a chromosome. Besides the functional genes and their alleles, there are also parts of natural genomes which have no (obvious) function [768, 769]. The American biochemist Gilbert [770] coined the term *intron*⁸ for such parts. Similar structures can also be observed in evolutionary algorithms.

⁴ You can find an illustration of the DNA in Figure 1.14 on page 24

⁵ http://en.wikipedia.org/wiki/Gray_coding [accessed 2007-07-03]

⁶ http://en.wikipedia.org/wiki/Cell_division [accessed 2007-07-03]

⁷ Matter of fact, cell division will continue until the individual dies. However, this is not important here.

⁸ <http://en.wikipedia.org/wiki/Intron> [accessed 2007-07-05]

Definition 3.2 (Intron). Parts of a genotype $g \in \mathbb{G}$ that does not contribute to the phenotype $x = \text{gpm}(g)$ are referred to as introns.

Biological introns have often been thought of as junk DNA or “old code”, i.e., parts of the genome that were translated to proteins in evolutionary past, but now are not used anymore. Recent discoveries indicate that introns are maybe not as useless as initially assumed [771]. Instead, they seem to provide support for efficient splicing, for instance. The role of introns in genetic algorithms is as same as mysterious. They represent a form of redundancy – which is known to have possible as well as negative effects, as outlined in Section 1.4.4 on page 46.

Figure 3.2 updates Figure 1.15 on page 26 and illustrates the relations between the aforementioned entities in a bit string genome $\mathbb{G} = \mathbb{B}^4$ of the length 4, where two bits encode for one coordinate in a two-dimensional plane. If additional bits are appended to the genotype – these occur as introns and will not influence the phenotype.




Fig. 3.2: A four bit string genome $\mathbb{G}$ and a fictitious phenotype $\mathbb{X}$.

3.4 Fixed-Length String Chromosomes

Especially widespread in genetic algorithms are search spaces based on fixed-length chromosomes. A lot of research has been conducted investigating the properties of their crossover and mutation operations [772, 766, 73].

3.4.1 Creation

Creation of fixed-length string individuals means simple to create a new tuple of the structure defined by the genome and initialize it randomized.

$$\text{create}_H() \equiv (g[1], g[2], \dots, g[n]) : g[i] = \mathbb{G}_i[\lfloor \text{random}_u() * \text{len}(\mathbb{G}_i) \rfloor] \quad \forall i \in 1..n \quad (3.3)$$

3.4.2 Mutation

Mutation is an important method of preserving the diversity of the solution candidates. In fixed-length string chromosomes, it can be achieved by modifying the value of one element of the genotype, as illustrated in Figure 3.3. More generally, a mutation may change $0 \leq n < |g|$ locations in the string. In binary coded chromosomes, for example, the elements are bits which are simply toggled. For real-encoded genomes, modifying an element g_i can be done by replacing it with a number drawn from a normal distribution with expected value g_1 , like $g_i^{\text{new}} \sim N(g_1, \sigma^2)$.





Fig. 3.3: Value-altering mutation of string chromosomes.

3.4.3 Permutation

The permutation operation is an alternative mutation method where the alleles of two genes are exchanged. This, of course, makes only sense if all genes have similar data types. Permutation is, for instance, useful when solving problems that involve finding an optimal sequence of items, like the traveling salesman problem. Here, a genotype g could encode the sequence in which the cities are visited. Exchanging two alleles then equals of switching two cities in the route.



Exchange the alleles of two genes

Fig. 3.4: Permutation applied to a string chromosome.

3.4.4 Crossover

Amongst all evolutionary algorithms, genetic algorithms have a recombination operation which is closest to the natural paragon. Figure 3.5 outlines the recombination of two string chromosomes, the so-called *crossover*, which is performed by swapping parts of two genotypes.

When performing crossover, both parental chromosomes are split at a randomly determined *crossover point*. Subsequently, a new child genotype is created by appending the first part of the first parent with the second part of the second parent. This method is called single-point crossover (SPX). In two-point crossover (TPX), both parental genotypes are split at two points, constructing a new offspring by using parts number one and three from the first, and the middle part from the second ancestor. The generalized form of this technique is the n -point crossover, also called multi-cpoint crossover (MPX). For fixed-length strings, the crossover points for both parents are always identical.




Fig. 3.5: Crossover (recombination) of fixed-length string chromosomes.

3.5 Variable-Length String Chromosomes

3.5.1 Creation

Variable-length strings can be created by first randomly drawing a length $l > 0$ and then creating a list of that length filled with random elements.

3.5.2 Mutation

If the string chromosomes are of variable length, the set of mutation operations introduced in Section 3.4 can be extended by two additional methods. On one hand, we could insert a couple of elements at any given position into a chromosome. On the other hand, this operation can be reversed by deleting elements from the string. These operations are one of the reasons why variable-length strings need to be constructed of elements of the same type. There is no longer a constant relation between locus and type. It should be noted that both, insertion and deletion, are also implicitly performed by crossover. Recombining two identical strings with each other can, for example, lead to deletion of genes. The crossover of different strings may turn out as an insertion of new genes into an individual.




Fig. 3.6: Mutation of variable-length string chromosomes.

3.5.3 Crossover

For variable-length string chromosomes, the same crossover operations are available as for fixed-length strings except that the strings are no longer necessarily split at the same loci. The length of the new strings resulting from such a *cut and splice* operation may differ from the length of the parents, as sketched in Figure 3.7. A special case of this recombination operation is the homologous crossover where only genes at the same loci are exchanged. You can find this method discussed in Section 4.6.1 on page 163.




Fig. 3.7: Crossover of variable-length string chromosomes.

3.6 Genotype-Phenotype Mappings and Artificial Embryogeny

As already stated a dozen times by now, genetic algorithms use string genomes $\mathbb{G}$ to encode the phenotypes x that represent the possible solutions. These phenotypes, however, do not necessarily need to be one-dimensional strings too. Instead, they can be construction plans, circuit layouts, or trees⁹. The process of translating genotypes into corresponding phenotypes is called genotype-phenotype mapping and has been introduced in Definition 1.30 on page 25.

Embryogenesis is the natural process in which the embryo forms and develops¹⁰ and to which the genotype-phenotype mapping in genetic algorithms and Genetic Programming corresponds. Most of even the more sophisticated of these mappings are based on an implicit one-to-one relation in terms of complexity. In the Grammar-guided Genetic Programming approach Gads¹¹, for example, a single gene encodes (at most) the application of a single grammatical rule, which in turn unfolds a single node in a tree.

Embryogeny in nature is much more complex. Among other things, the DNA, for instance, encodes the structural design information of the human brain. As pointed out by Manos et al. [773], there are only about 30 thousand active genes in the human genome (2800 million amino acids) for over 100 trillion neural connections in our cerebrum. A huge manifold of information is hence decoded from “data” which is of a much lower magnitude. This is possible because the same genes can be reused in order to repeatedly create the same pattern. The layout of the light receptors in the eye, for example, is always the same – just their wiring changes.

Definition 3.3 (Artificial Embryogeny). We subsume all methods of transforming a genotype into a phenotype of (much) higher complexity under the subject of *artificial embryogeny* [773, 774, 775] (also known as computational embryogeny [776, 777]).

Two different approaches are common in artificial embryogeny: constructing the phenotype by using a grammar to translate the genotype and expanding it step by step until a terminal state is reached or simulating chemical processes. Both methods may also require subsequent correction steps that ensure that the produced results are correct, which is also common in normal genotype-phenotype mappings [778]. An example for gene reuse is the genotype-phenotype mapping performed in Grammatical Evolution which is discussed in Section 4.5.6 on page 152.

3.7 Schema Theorem

*“Now if you were to alter masks every time fame circus approaches
Do you really think your maker wouldn’t notice?”*

— Aesop Rock [779]

⁹ See for example Section 4.5.6 on page 151

¹⁰ <http://en.wikipedia.org/wiki/Embryogenesis> [accessed 2007-07-03]

¹¹ See Section 4.5.5 on page 149 for more details.

The Schema Theorem is a special instance of forma analysis (discussed in Section 1.5.1 on page 56) in genetic algorithms. Matter of fact, it is older than its generalization and was first stated by Holland back in 1975 [253, 254, 80].

3.7.1 Schemata and Masks

Assume that the genotypes g in the search space $\mathbb{G}$ of genetic algorithms are strings of a fixed-length l over an alphabet¹² Σ , i. e., $\mathbb{G} = \Sigma^l$. Normally, Σ is the binary alphabet $\Sigma = \{\text{true}, \text{false}\} = \{0, 1\}$. From forma analysis, we know that properties can be defined on the genotypic or the phenotypic space. For fixed-length string genomes, we can consider the values at certain loci as properties of a genotype. There are two basic principles on defining such properties: masks and do not care symbols.

Definition 3.4 (Mask). For a fixed-length string genome $\mathbb{G} = \Sigma^l$, we define the set of all genotypic masks M_l as the power set¹³ of the valid loci $Ml = \mathcal{P}(\{1, \dots, l\})$ [255]. Every mask $m_i \in M_l$ defines a property ϕ_i and an equivalence relation:

$$g \sim_{\phi_i} h \Leftrightarrow g[j] = h[j] \quad \forall j \in m_i \quad (3.4)$$

The order “order(m_i)” of the mask m_i is the number of loci defined by it:

$$\text{order}(m_i) = |m_i| \quad (3.5)$$

The defined length $\delta(m_i)$ of a mask m_i is the maximum distance between two indices in the mask:

$$\delta(m_i) = \max\{|j - k| \mid \forall j, k \in m_i\} \quad (3.6)$$

A mask contains the indices of all elements in a string that are interesting in terms of the property it defines. Assume we have bit strings of the length $l = 3$ as genotypes ($\mathbb{G} = \mathbb{B}^3$). The set of valid masks M_3 is then $M_3 = \{\{1\}, \{2\}, \{3\}, \{1, 3\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$. The mask $m_1 = \{1, 2\}$, for example, specifies that the values at the loci 1 and 2 of a genotype denote the value of a property ϕ_1 and the value of the bit at position 3 is irrelevant. Therefore, it defines four formae $A_{\phi_1=(0,0)} = \{(0, 0, 0), (0, 0, 1)\}$, $A_{\phi_1=(0,1)} = \{(0, 1, 0), (0, 1, 1)\}$, $A_{\phi_1=(1,0)} = \{(1, 0, 0), (1, 0, 1)\}$, and $A_{\phi_1=(1,1)} = \{(1, 1, 0), (1, 1, 1)\}$.

Definition 3.5 (Schema). A forma defined on a string genome concerning the values of the characters at specified loci is called *Schema* [253, 448].

3.7.2 Wildcards

The second method of specifying such schemata is to use *don't care* symbols (wildcards) to create “blueprints” H of their member individuals. Therefore, we place the don't care symbol $*$ at all irrelevant positions and the characterizing values of the property at the others.

$$\forall j \in 1..l \Rightarrow H[j] = \begin{cases} g[j] & \text{if } j \in m_i \\ * & \text{otherwise} \end{cases} \quad (3.7)$$

$$H[j] \in \Sigma \cup \{*\} \quad \forall j \in 1..l \quad (3.8)$$

$$(3.9)$$

We now can redefine the aforementioned schemata like: $A_{\phi_1=(0,0)} \equiv H1 = (0, 0, *)$, $A_{\phi_1=(0,1)} \equiv H1 = (0, 1, *)$, $A_{\phi_1=(1,0)} \equiv H1 = (1, 0, *)$, and $A_{\phi_1=(1,1)} \equiv H1 = (1, 1, *)$.

These schemata mark hyperplanes in the search space $\mathbb{G}$, as illustrated in Figure 3.8 for the three bit genome.

¹² Alphabets and such and such are defined in Section 29.3 on page 501.

¹³ The power set you can find described in Definition 26.9 on page 410.




Fig. 3.8: An example for schemata in a three bit genome.

3.7.3 Holland's Schema Theorem

The Schema Theorem¹⁴ was defined by Holland [253] for genetic algorithms which use fitness-proportionate selection (see Section 2.4.3 on page 95) and where fitness is subject to maximization [254, 80].

$$\text{countItemOccurrences}(H, \text{Pop})_{t+1} \geq \frac{\text{countItemOccurrences}(H, \text{Pop})_t * \bar{v}(H)_t}{\bar{v}_t} \quad (3.10)$$

where

- $\text{countItemOccurrences}(H, \text{Pop})_t$ is the number of instances of a given schema defined by the blueprint H in the population Pop of generation t ,
- $\bar{v}(H)_t$ is the average fitness of the members of this schema (observed in time step t),
- $\bar{v}_t$ is the average fitness of the population in time step t , and
- p is the probability that an instance of the schema will be “destroyed” by a reproduction operation, i. e., the probability that the offspring of an instance of the schema is not an instance of the schema.

From this formula can be deduced that genetic algorithms will generate for short, above-average fit schemata an exponentially rising number of samples. This is because they will multiply with a certain factor in each generation and only few of them are destroyed by the reproduction operations.

3.7.4 Criticism of the Schema Theorem

The deduction that good schemata will spread exponentially is only a very optimistic assumption and not generally true. If a highly fit schema has many offspring with good fitness, this will also improve the overall fitness of the population. Hence, the probabilities in Equation 3.10 will shift over time. Generally, the Schema Theorem represents a lower bound that will only hold for one generation [645]. Trying to derive predictions for more than one or two generations using the Schema Theorem as is will lead to deceptive or wrong results [417, 780].

¹⁴ http://en.wikipedia.org/wiki/Holland%27s_Schema_Theorem [accessed 2007-07-29]

Furthermore, the population of a genetic algorithm only represents a sample of limited size of the search space $\mathbb{G}$. This limits the reproduction of the schemata but also makes statements about probabilities in general more complicated. Since we only have samples of the schemata H and cannot be sure if $\bar{v}(H)_t$ really represents the average fitness of all the members of the schema (that's why we annotate it with $_t$ instead of writing $\bar{v}(H)$). Thus, even reproduction operators that preserve the instances of the schema may lead to a decrease of $\bar{v}(H)_{t+..}$ by time. It is also possible that parts of the population already have converged and other members of a schema will not be explored anymore, so we do not get further information about its real utility.

Additionally, we cannot know if it is really good if one specific schema spreads fast, even if it is very fit. Remember that we have already discussed the exploration versus exploitation topic and the importance of diversity in Section 1.4.1 on page 39.

Another issue is that we implicitly assume that most schemata are compatible and can be combined, i. e., that there is low interaction between different genes. This is also not generally valid. The expressiveness of masks and blueprints even is limited and can be argued that there are properties which we cannot specify with them. Take the set D_3 of numbers divisible by three for example $D_3 = \{3, 6, 9, 12, \dots\}$. Representing them as binary strings will lead to $D_3 = \{0011, 0110, 1001, 1100, \dots\}$ if we have a bit-string genome of the length 4. Obviously, we cannot seize these genotypes in a schema using the discussed approach. They may, however, be gathered in a forma. The Schema Theorem, however, cannot hold for such a forma since the probability p of destruction may be different from instance to instance.

3.7.5 The Building Block Hypothesis

The Building Block Hypothesis (BBH) proposed by Goldberg [139], Holland [253] is based on two assumptions:

1. When a genetic algorithm solves a problem, there exist some low-order, low-defining length schemata with above-average fitness (the so-called *building blocks*).
2. These schemata are combined step by step by the genetic algorithm in order to form larger and better strings. By using the building blocks instead of testing any possible binary configuration, genetic algorithms efficiently decrease the complexity of the problem. [139]

Although it seems as if the Building Block Hypothesis was supported by the Schema Theorem, this cannot be verified easily. Experiments that originally were intended to proof this theory often did not work out as planned [781] (and also consider the criticisms of the Schema Theorem mentioned in the previous section). In general, there exists much criticism of the Building Block Hypothesis and, although it is a very nice train of thought, it can be regarded as not proven sufficiently.

Genetic Programming

4.1 Introduction

The term *Genetic Programming*¹ [65, 1] has two possible meanings. On one hand, it is often used to subsume all evolutionary algorithms that produce tree data structures as phenotypes. On the other hand, we can also define it as the set of all evolutionary algorithms that breed programs², algorithms, and similar constructs. In this chapter, we focus on the latter definition, which still includes discussing tree-shaped genomes.

The conventional well-known input-processing-output model³ of information processing in computer science states that a running instance of a program uses its input information to compute and return output data. In Genetic Programming, we already know (or are able to produce) some inputs or situations and corresponding output data samples. What we want to find is a program that connects them or that exhibits some kind of desired behavior according to the specified situations, as sketched in Figure 4.1.




Fig. 4.1: Genetic Programming in the context of the IPO model.

4.1.1 History

The history of Genetic Programming [782] goes back to the early days of computer science. In 1957, Friedberg [783] left the first footprints in this area by using a learning algorithm to stepwise improve a program. The program was represented as a sequence of instructions⁴ for a theoretical computer called *Herman* [783, 784]. Friedberg did not use an evolutionary, population-based approach for searching the programs. This may be because the idea of

¹ http://en.wikipedia.org/wiki/Genetic_programming [accessed 2007-07-03]

² We have extensively discussed the topic of algorithms and programs in Section 29.1.1 on page 479.

³ see Section 29.1.1 on page 481

⁴ Linear Genetic Programming is discussed in Section 4.6 on page 161.

evolutionary algorithms wasn't fully developed yet⁵ and also because of the limited computational capacity of the computers of that era.

The evolutionary programming approach for evolving finite state machines by Fogel et al. [785], discussed in Chapter 6 on page 195, dates back to 1966. In order to build predictive behaviors, different forms of mutation (but no crossover) were used for creating offspring from successful individuals.

Fourteen years later, the next generation of scientists began to look for ways to evolve programs. New results were reported by Smith [786] in his PhD thesis in 1980. Forsyth [787] evolved trees denoting fully bracketed Boolean expressions for classification problems in 1981 [787, 788, 789].

The mid-1980s were a very productive period for the development of Genetic Programming. Cramer [790] applied a genetic algorithm in order to evolve a program written in a subset of the programming language PL in 1985.⁶ This GA used a string of integers as genome and employed a genotype-phenotype mapping that recursively transformed them into program trees. At the same time, the undergraduate student Schmidhuber [791] also used a genetic algorithm to evolve programs at the Siemens AG. He re-implemented his approach in *Prolog* at the TU Munich in 1987 [792, 791]. Hicklin [793] and Fujiki [794] implemented reproduction operations for manipulating the if-then clauses of LISP programs consisting of single COND-statements. With this approach, Fujiko and Dickinson [795] evolved strategies for playing the iterated prisoner's dilemma game. Bickel and Bickel [796] evolved sets of rules which were represented as trees using tree-based mutation crossover operators.

Genetic Programming became fully accepted at the end of this productive decade mainly because of the work of Koza [797; 798]. He also studied many benchmark applications of Genetic Programming, such as learning of Boolean functions [799, 800], the Artificial Ant problem⁷ [801, 802, 65], and symbolic regression⁸ [799, 65], a method for obtaining mathematical expressions that match given data samples. Koza formalized (and patented [797, 803]) the idea of employing genomes purely based on tree data structures rather than string chromosomes as used in genetic algorithms. In symbolic regression, such trees can, for instance, encode *Lisp* S-expressions⁹ where a node stands for a mathematical operation and its child nodes are the parameters of the operation. Leaf nodes then are terminal symbols like numbers or variables. This form of Genetic Programming is called *Standard Genetic Programming* or SGP, in short. With it, not only mathematical functions but also more complex programs can be expressed as well.

Generally, a tree can represent a rule set [804, 805], a mathematical expressions, a decision tree [806], or even the blueprint of an electrical circuit [807]. Trees are very close to the natural structure of algorithms and programs. The syntax of most of the high-level programming languages, for example, leads to a certain hierarchy of modules and alternatives. Not only does this form normally constitute a tree – compilers even use tree representations internally. When reading the source code of a program, they first split it into tokens¹⁰, parse¹¹ these tokens, and finally create an abstract syntax tree¹² (AST) [808, 809]. The internal nodes of ASTs are labeled by operators and the leaf nodes contain the operands of these operators. In principle, we can illustrate almost every¹³ program or algorithm as such an AST (see Figure 4.2).

⁵ Compare with Section 3.1 on page 113.

⁶ Cramer's approach is discussed in Section 4.4.1 on page 141.

⁷ The Artificial Ant is discussed in Section 20.3.1 on page 303 in this book.

⁸ More information on symbolic regression is presented in Section 22.1 on page 341 in this book.

⁹ List S-expressions are discussed in Section 29.3.11 on page 512

¹⁰ http://en.wikipedia.org/wiki/Lexical_analysis [accessed 2007-07-03]

¹¹ http://en.wikipedia.org/wiki/Parse_tree [accessed 2007-07-03]

¹² http://en.wikipedia.org/wiki/Abstract_syntax_tree [accessed 2007-07-03]

¹³ Excluding such algorithms and programs that contain jumps (the infamous “goto”) that would produce crossing lines in the flowchart (<http://en.wikipedia.org/wiki/Flowchart> [accessed 2007-07-03]).




Fig. 4.2: The AST representation of algorithms/programs.

Tree-based Genetic Programming directly evolves individuals in this form, which also provides a very intuitive representation for mathematical functions for which it has initially been used for by Koza. Another interesting aspect of the tree genome is that it has no natural role model. While genetic algorithms match their direct biological metaphor particularly well, Genetic Programming introduces completely new characteristics and traits. Genetic Programming is one of the few techniques that are able to learn solutions of potentially unbound complexity. It can be considered as more general than genetic algorithms, because it makes fewer assumptions about the structure of possible solutions. Furthermore, it often offers white-box solutions that are human-interpretable. Other optimization approaches like artificial neural networks, for example, generate black-box outputs, which are highly complicated if not impossible to fully grasp [810].

4.2 General Information

4.2.1 Areas Of Application

Some example areas of application of Genetic Programming are:

Application	References
Symbolic Regression and Function Synthesis	[799, 65, 811, 812, 813, 65, 814, 815] Section 22.1
Grammar Induction	[816, 817, 818, 819]
Data Mining and Data Analysis	[820, 821, 822, 823, 824, 464, 806, 825] Section 21.1.2
Electrical Engineering and Circuit Design	[807, 826, 827, 828, 829, 830, 831]
Medicine	[832, 833, 834, 835]
Economics and Finance	[836, 837, 838, 839]
Geometry and Physics	[840, 841]
Cellular Automata and Finite State Machines	[842, 843, 844, 845]
Automated Programming	[846, 57, 847, 848, 849, 850]
Robotics	[851, 852, 853, 854, 849, 855, 856, 854, 857] [858, 859, 860, 861, 862, 863, 864, 865, 866, 867, 868, 869, 870, 871, 872, 873] Section 23.2 on page 365
Networking and Communication	[802, 874, 875, 876, 877, 878, 879, 880, 641, 881, 882, 883, 884, 885, 886, 887, 854, 888, 889, 890, 891, 857]
Evolving Behaviors, e.g., for Agents or Game Players	[892, 893, 894, 895, 896]
Pattern Recognition	[897, 898]
Biochemistry	[899, 900]
Machine Learning	

See also Section 4.4.3 on page 144, Section 4.5.6 on page 154, and Section 4.7.4 on page 169.

4.2.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on Genetic Programming are:

EuroGP: European Conference on Genetic Programming

<http://www.evostar.org/> [accessed 2007-09-05]

Co-located with EvoWorkshops and EvoCOP.

History: 2008: Naples, Italy, see [901]
 2007: Valencia, Spain, see [902]
 2006: Budapest, Hungary, see [903]
 2005: Lausanne, Switzerland, see [904]
 2004: Coimbra, Portugal, see [905]
 2003: Essex, UK, see [906]
 2002: Kinsale, Ireland, see [907]
 2001: Lake Como, Italy, see [908]
 2000: Edinburgh, Scotland, UK, see [909]
 1999: Göteborg, Sweden, see [910]

1998: Paris, France, see [911, 912]

GECCO: Genetic and Evolutionary Computation Conference

see Section 2.2.2 on page 79

GP: Annual Genetic Programming Conference

Now part of GECCO, see Section 2.2.2 on page 79

History: 1998: Madison, Wisconsin, USA, see [913, 914]

1997: Stanford University, CA, USA, see [915, 916]

1996: Stanford University, CA, USA, see [917, 918]

GPTP: Genetic Programming Theory Practice Workshop

<http://www.cscs.umich.edu/gptp-workshops/> [accessed 2007-09-28]

History: 2007: Ann Arbor, Michigan, USA, see [919]

2006: Ann Arbor, Michigan, USA, see [920]

2005: Ann Arbor, Michigan, USA, see [921]

2004: Ann Arbor, Michigan, USA, see [922]

2003: Ann Arbor, Michigan, USA, see [923]

ICANNGA: International Conference on Adaptive and Natural Computing Algorithms

see Section 2.2.2 on page 79

Mendel: International Conference on Soft Computing

see Section 1.6.2 on page 63

4.2.3 Journals

Some journals that deal (at least partially) with Genetic Programming are:

Genetic Programming and Evolvable Machines (GPEM), ISSN: 1389-2576 (Print) 1573-7632 (Online), appears quarterly, editor(s): Wolfgang Banzhaf, publisher: Springer Netherlands, <http://springerlink.metapress.com/content/104755/> [accessed 2007-09-28]

4.2.4 Online Resources

Some general, online available resources on Genetic Programming are:

<http://www.genetic-programming.org/> [accessed 2007-09-20] and <http://www.genetic-programming.com/> [accessed 2007-09-20]

Last update: up-to-date

Description: Two portal pages on Genetic Programming websites, both maintained by Koza.

<http://www.cs.bham.ac.uk/~wbl/biblio/> [accessed 2007-09-16]

Last update: up-to-date

Description: Langdon's large Genetic Programming bibliography.

http://www.lulu.com/items/volume_63/2167000/2167025/2/print/book.pdf [accessed 2008-03-26]

Last update: up-to-date

Description: A Field Guide to Genetic Programming, see [924]

<http://www.aaii.org/AITopics/html/genalg.html> [accessed 2008-05-17]

Last update: up-to-date

Description: The genetic algorithms and Genetic Programming pages of the AAI

http://www.cs.ucl.ac.uk/staff/W.Langdon/www_links.html [accessed 2008-05-18]

Last update: 2007-07-28

Description: William Langdon's Genetic Programming contacts

4.2.5 Books

Some books about (or including significant information about) Genetic Programming are:

-
- Koza [65]: *Genetic Programming, On the Programming of Computers by Means of Natural Selection*
- Poli, Langdon, and McPhee [924]: *A Field Guide to Genetic Programming*
- Koza [925]: *Genetic Programming II: Automatic Discovery of Reusable Programs: Automatic Discovery of Reusable Programs*
- Koza, Bennett III, Andre, and Keane [926]: *Genetic Programming III: Darwinian Invention and Problem Solving*
- Koza, Keane, Streeter, Mydlowec, Yu, and Lanza [850]: *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*
- Langdon and Poli [57]: *Foundations of Genetic Programming*
- Langdon [927]: *Genetic Programming and Data Structures: Genetic Programming + Data Structures = Automatic Programming!*
- Banzhaf, Nordin, Keller, and Francone [846]: *Genetic Programming: An Introduction – On the Automatic Evolution of Computer Programs and Its Applications*
- Kinnear, Jr. [928]: *Advances in Genetic Programming, Volume 1*
- Angeline and Kinnear, Jr [929]: *Advances in Genetic Programming, Volume 2*
- Spector, Langdon, O'Reilly, and Angeline [930]: *Advances in Genetic Programming, Volume 3*
- Brameier and Banzhaf [931]: *Linear Genetic Programming*
- Wong and Leung [825]: *Data Mining Using Grammar Based Genetic Programming and Applications*
- Geyer-Schulz [932]: *Fuzzy Rule-Based Expert Systems and Genetic Machine Learning*
- Spector [933]: *Automatic Quantum Computer Programming – A Genetic Programming Approach*
-

4.3 (Standard) Tree Genomes


Tree-based Genetic Programming (TGP) is, not alone for historical reasons, the most widespread Genetic Programming variant. In this section, we introduce the well-known reproduction operations applicable to tree genomes.

4.3.1 Creation


Before the evolutionary process can begin, we need an initial, randomized population. In genetic algorithms, we therefore simply created a set of random bit strings. For Genetic Programming, we do the same with trees instead of such one-dimensional sequences.

Normally (but not necessarily), there is a maximum depth d specified that the tree individuals are not allowed to surpass. Thus, the creation operation will return only trees where the path between the root and the most distant leaf node is not longer than d . There are three different ways for realizing the “create()” operation (see Definition 2.9 on page 108) for trees which can be distinguished according to the depth of the produced individuals.

The *full* method (Figure 4.3) creates trees where each non-backtracking path from the root to the leaf nodes has exactly the length d .

Fig. 4.3: Tree creation by the *full* method.

The *grow* method depicted in Figure 4.4, on the other hand, creates trees where each non-backtracking path from the root to the leaf nodes is not longer than d but may be shorter. This is achieved by deciding randomly for each node if it should be a leaf or not when it is attached to the tree. Of course, to nodes of the depth $d - 1$, only leaf nodes may be attached to.

Fig. 4.4: Tree creation by the *grow* method.

Koza additionally introduced a mixture method called *ramped half-and-half* [65]. For each tree to be created, this algorithm draws a number r uniformly distributed between 2 and d : ($r = \lfloor \text{random}2d + 1 \rfloor$). Now either *full* or *grow* is chosen to finally create a tree with the maximum depth r (instead of d). This method is preferable since it produces an especially wide range of different tree depths and shapes and thus provides a great initial diversity.

4.3.2 Mutation

Analogously to nature, tree genotypes may undergo small variations called mutations during the reproduction process. Mutation of a tree is defined as randomly selecting a node, removing this node and all its children and replacing it with a new, randomly created one [65]. This general definition subsumes

- insertions of new nodes or small trees,
- replacement of existing nodes with others, or
- the deletion of nodes, as illustrated in Figure 4.5.

4.3.3 Recombination

The mating process in nature, the recombination of the genotypes of two individuals, also takes place in tree-based Genetic Programming. Applying the default sub-tree exchange recombination operator to two trees means to swap sub-trees between them as illustrated




Fig. 4.5: Possible tree mutation operations.




Fig. 4.6: Tree crossover by exchanging sub-trees.

in Figure 4.6. Therefore, one single sub-tree is selected randomly from each of the parents. They are cut out and reinserted in the partner genotype.

Notice that, like in genetic algorithms, the insertion and deletion of nodes can also be achieved by recombination. If a depth restriction is imposed on the genome, both, the mutation and the crossover operation have to respect them. The new trees they create must not exceed it.

The intent of using the recombination operation in Genetic Programming is the same as in genetic algorithms. Over many generations, successful building blocks – for example a highly fit expression in a mathematical formula – should spread throughout the population and be combined with good genes of different solution candidates. Yet, recombination in Standard Genetic Programming can also have a very destructive effect on the individual fitness [260, 934, 846]. Angeline [935] even argues that it performs no better than mutation and causes bloat [936].

Several techniques have been proposed in order to mitigate these effects [937]. In 1994, D’Haeseleer [938] obtained modest improvements with his strong context preserving crossover that permitted only the exchange of sub-trees that occupied the same positions in the parents. Poli and Langdon [939; 940] define the similar single-point crossover for the same purpose: increasing the probability of exchanging genetic material which is struc-

tural and functional akin and thus decreasing the disruptiveness. A related approach define by Francone et al. [937] for linear Genetic Programming is discussed in Section 4.6.1 on page 163.

4.3.4 Permutation

The tree permutation operation illustrated in Figure 4.7 resembles the permutation operation of string genomes. Like mutation, it is used to reproduce one single tree asexually. It first selects an internal node of the parental tree. The child nodes attached to that node are then shuffled randomly, i. e., permuted. If the tree represents a mathematical formula and the operation represented by the node is commutative, this has no direct effect. The main goal is to re-arrange the nodes in highly fit sub-trees in order to make them less fragile for other operations such as recombination. The effects of this operation are doubttable and most often it is not applied [65].




Fig. 4.7: Tree permutation – asexually shuffling sub-trees.

4.3.5 Editing

Editing trees in Genetic Programming is what simplifying is to mathematical formulas. Take $x = b + (7 - 4) + (1 * a)$ for instance. This expression clearly can be written in a shorter way by replacing $(7 - 4)$ with 3 and $(1 * a)$ with a . By doing so, we improve its readability and also decrease the computational time for concrete values of a and b . Similar measures can often be applied to algorithms and program code. Editing a tree as outlined in Figure 4.8 means to create a new offspring tree which is more efficient but, in terms of functional aspects, equivalent to its parent. It is thus a very domain-specific operation.




Fig. 4.8: Tree editing – asexual optimization.

A positive aspect of editing is that it usually reduces the number of nodes in a tree by removing useless expression, for instance. This makes it more easy for recombination operations to pick “important” building blocks. At the same time, the expression $(7 - 4)$ is now less likely to be destroyed by the reproduction processes since it is replaced by the single terminal node 3.

On the other hand, editing also reduces the diversity in the genome which could degrade the performance by decreasing the variety of structures available. Another negative aspect would be if (in our example) a fitter expression was $(7 - (4 * a))$ and a is a variable close to 1. Then, transforming $(7 - 4)$ into 3 prevents a transition to the fitter expression.

In Koza's experiments, Genetic Programming with and without editing showed equal performance, so this operation is not necessarily needed [65].

4.3.6 Encapsulation

The idea behind the encapsulation operation is to identify potentially useful sub-trees and to turn them into atomic building block as sketched in Figure 4.9. To put it plain, we create new terminal symbols that (internally hidden) are trees with multiple nodes. This way, they will no longer be subject to potential damage by other reproduction operations. The new terminal may spread throughout the population in the further course of the evolution. According to Koza, this operation has no substantial effect but may be useful in special applications like the evolution of artificial neural networks [65].




Fig. 4.9: An example for tree encapsulation.

4.3.7 Wrapping

Applying the wrapping operation means to first select an arbitrary node n in the tree. Additionally, we create a new non-terminal node m outside of the tree. In m , at least one child node position is left unoccupied. We then cut n (and all its potential child nodes) from the original tree and append it to m by plugging it into the free spot. Now we hang m into the tree position that formerly was occupied by n .




Fig. 4.10: An example for tree wrapping.

The purpose of this reproduction method illustrated in Figure 4.10 is to allow modifications of non-terminal nodes that have a high probability of being useful. Simple mutation would, for example, cut n from the tree or replace it with another expression.

This will always change the meaning of the whole sub-tree below n dramatically, like for example in $(b+3) + a \rightarrow (b*3) + a$. By wrapping however, a more subtle change like $(b+3) + a \rightarrow ((b+1)+3) + a$ is possible.

The wrapping operation is introduced by the author – at least, I have not seen another source where it is used.

4.3.8 Lifting

While wrapping allows nodes to be inserted in non-terminal positions with small change of the tree's semantic, lifting is able to remove them in the same way. It is the inverse operation to wrapping, which becomes obvious when comparing Figure 4.10 and Figure 4.11.




Fig. 4.11: An example for tree lifting.

Lifting begins with selecting an arbitrary inner node n of the tree. This node then replaces its parent node. The parent node inclusively all of its child nodes (except n) are removed from the tree. With lifting, a tree that represents the mathematical formula $(b + (1 - a)) * 3$ can be transformed to $b * 3$ in a single step. Lifting is used by the author in his experiments with Genetic Programming (see for example Section 23.2.2 on page 366). I, however, have not yet found other sources using a similar operation.

4.3.9 Automatically Defined Functions

Automatically defined functions (ADFs) introduced by Koza [65] provide some sort of pre-specified modularity for Genetic Programming. Finding a way to evolve modules and reusable building blocks is one of the key issues in using GP to derive higher-level abstractions and solutions to more complex problems [883, 925]. If ADFs are used, a certain structure is defined for the genome. The root of the trees loses its functional responsibility and now serves only as glue that holds the individual together. It has a fixed number n of children, from which $n - 1$ are automatically defined functions and one is the result-generating branch. When evaluating the fitness of the individual, only this result-generating branch is taken into consideration whereas the root and the ADFs are ignored. The result-generating branch, however, may use any of the automatically defined functions to produce its output.

When ADFs are employed, not only their number must be specified beforehand but also the number of arguments of each of them. How this works can maybe best illustrated using an example from function approximation¹⁴, since this is the area where the idea originally stems from. In Fig. 4.12.2 we illustrate such a concrete example while Fig. 4.12.1 outlines how a genotype with s could look like.

In the example, we want to approximate a function g with the one parameter x . We use a genome where two functions (let us call them f_0 and f_1) are automatically defined. f_0 has a single formal parameter a and f_1 has two formal parameters a and b . The genotype Fig. 4.12.2 encodes the following mathematical functions:

¹⁴ A very common example for function approximation, Genetic Programming-based symbolic regression, is discussed in Section 22.1 on page 341.




Fig. 4.12.1: general structure




Fig. 4.12.2: example

Fig. 4.12: Automatically defined functions in Genetic Programming.

$$\begin{aligned}
 f_0(a) &= a + 7 \\
 f_1(a, b) &= (-a) * b \\
 g &\approx f_1(4, f_0(x)) + (f_0(x) * 3)
 \end{aligned}$$

Hence, $g \approx ((-4) * (x+7)) + (x+7) * 3$. As you can see, the number of children of the function calls in the result-generating branch must be equal to the number of the parameters of the corresponding ADF.

Although ADFs were first introduced in symbolic regression by Koza [65], they can also be applied to a variety of other problems like in the evolution of agent behaviors [876, 877, 941, 641], electrical circuit design [828], or the evolution of robotic behavior [855].

4.3.10 Automatically Defined Macros

Spector's idea of automatically defined macros (ADMs) complements the ADFs of Koza [942, 943]. Both concepts are very similar and only differ in the way that their parameters are handled. The parameters in automatically defined functions are always values whereas automatically defined macros work on code expressions. This difference shows up only when side-effects come into play.

In Figure 4.13, we have illustrated the pseudo-code of two programs – one with a function (called ADF) and one with a macro (called ADM). Each program has a variable x which is initially zero. The function $y()$ has the side-effect that it increments x and returns its new value. Both, the function and the macro, return a sum containing their parameter a two times. The parameter of ADF is evaluated *before* ADF is invoked. Hence, x is incremented one time and 1 is passed to ADF which then returns $2=1+1$. The parameter of the macro, however, is the invocation of $y()$, not its result. Therefore, the ADM resembles to two calls to $y()$, resulting in x being incremented two times and in $3=1+2$ being returned.

In summary, automatically defined macros and automatically defined functions are very similar. Automatically defined macros are likely to be useful in scenarios where context-sensitive or side-effect-producing operators play important roles [942, 943]. In other scenarios, there is no much difference between the application of ADFs and ADMs. Finally, it should be mentioned that the concepts of automatically defined functions and macros are not restricted to the standard tree genomes but are also applicable in other forms of Genetic Programming, such as linear Genetic Programming or PADO.¹⁵

¹⁵ Linear Genetic Programming is discussed in Section 4.6 on page 161 and a summary on PADO can be found in Section 4.7.1 on page 163.




Fig. 4.13: Comparison of functions and macros.

4.3.11 Node Selection

In most of the reproduction operations for tree genomes, in mutation as well as in recombination, certain nodes in the trees need to be selected. In order to apply the mutation, we first need to find the node which is to be altered. For recombination, we need one node in each parent tree. These nodes are then exchanged. The question how to select these nodes seems to be more or less irrelevant but plays an important role in reality. The literature most often speaks of “randomly selecting” a node but does not describe how exactly this should be done.

A good method for doing so could select all nodes c and n in the tree t with exactly the same probability as done by the method “uniformSelectNode”, i. e., $P(\text{uniformSelectNode}(t) = c) = P(\text{uniformSelectNode}(t) = n) \forall c, n \in t$.

Therefore, we define the weight $\text{nodeWeight}(n)$ of a tree node n to be the total number of nodes in the sub-tree with n as root, i. e., itself, its children, grandchildren, grand-grandchildren, etc.

$$\text{nodeWeight}(n) = 1 + \sum_{i=0}^{\text{len}(n.\text{children})-1} \text{nodeWeight}(n.\text{children}[i]) \quad (4.1)$$

Thus, the `nodeWeight` of the root of a tree is the number of all nodes in the tree and the `nodeWeight` of each of the leaves is exactly 1. In `uniformSelectNode`, the probability for a node of being selected in a tree t is thus $1/\text{nodeWeight}(t)$. We can create such a probability distribution by descending it from the root according to Algorithm 4.1.

Algorithm 4.1: $n \leftarrow \text{uniformSelectNode}(t)$

Input: t : the (root of the) tree to select a node from
Data: c : the currently investigated node
Data: $c.\text{children}$: the list of child nodes of c
Data: b, d : two Boolean variables
Data: r : a value uniformly distributed in $[0, \text{nodeWeight}(c)]$
Data: i : an index
Output: n : the selected node

```

1 begin
2    $b \leftarrow \text{true}$ 
3    $c \leftarrow t$ 
4   while  $b$  do
5      $r \leftarrow \lfloor \text{random}_u(0, \text{nodeWeight}(c)) \rfloor$ 
6     if  $r \geq \text{nodeWeight}(c) - 1$  then  $b \leftarrow \text{false}$ 
7     else
8        $i \leftarrow \text{len}(c.\text{children}) - 1$ 
9       while  $i \geq 0$  do
10         $r \leftarrow r - \text{nodeWeight}(c.\text{children}[i])$ 
11        if  $r < 0$  then
12           $c \leftarrow c.\text{children}[i]$ 
13           $i \leftarrow -1$ 
14        else
15           $i \leftarrow i - 1$ 
16   return  $c$ 
17 end
```

A tree descend where with probabilities different from these defined here will lead to unbalanced node selection probability distributions. Then, the reproduction operators will prefer accessing some parts of the trees while very rarely altering the other regions. We could, for example, descend the tree by starting at the root t and would return the current node with probability 0.5 or recursively go to one of its children (also with 50% probability). Then, the root t would have a 50 : 50 chance of being the starting point of reproduction operation. Its direct children have at most probability $0.5^2/\text{len}(t.\text{children})$ each, and their children even $0.5^3/\text{len}(t.\text{children})\text{len}(t.\text{children}[i].\text{children})$ and so on. Hence, the leaves would almost never take actively part in reproduction. We could also choose other probabilities which strongly prefer going down to the children of the tree, but then, the nodes near to the root will most likely be left untouched during reproduction. Often, this approach is favored by selection methods, although leaves in different branches of the tree are not chosen with the same probabilities if the branches differ in depth. When applying Algorithm 4.1 on the other hand, there exist no regions in the trees that have lower selection probabilities than others.

4.4 Genotype-Phenotype Mappings

Genotype-phenotype mappings (GPM, see Section 3.6 on page 122) are used in many different Genetic Programming approaches. Here we give a few examples about them. Many of the Grammar-guided Genetic Programming approaches discussed in Section 4.5 on page 147 are based on similar mappings.

4.4.1 Cramer's Genetic Programming

It is interesting to see that the earliest Genetic Programming approaches were based on a genotype-phenotype mapping. One of them, dating back to 1985, is the method of Cramer [790]. His goal was to evolve programs in a modified subset of the programming language PL. Two simple examples for such programs, obtained from his work, are:

```

1  ;;Set variable V0 to have the value of V1
2  (:ZERO V0)
3  (:LOOP V1 (:INC V0))
4
5  ;;Multiply V3 by V4 and store the result in V5
6  (:ZERO V5)
7  (:LOOP V3 (:LOOP V4 (:INC V5)))

```

Listing 4.1: Two examples for the PL dialect used by Cramer for Genetic Programming

On basis of a genetic algorithm working on integer strings, he proposed two ideas on how to convert these strings to valid program trees.

The JB Mapping

The first approach was to divide the integer string into tuples of a fixed length which is large enough to hold the information required to encode an arbitrary instruction. In the case our examples, these are triplets where the first item identifies the operation, and the following two numbers define its parameters. Superfluous information, like a second parameter for a unary operation, is ignored.

```

1  (0 4 2) → (:BLOCK AS4 AS2)
2  (1 6 0) → (:LOOP V6 AS0)
3  (2 1 9) → (:SET V1 V9)
4  (3 17 8) → (:ZERO V17) ;;the 8 is ignored
5  (4 0 5) → (:INC V0) ;;the 5 is ignored

```

Listing 4.2: An example for the JB Mapping

Here, the symbols of the form Vn and ASn represent variables and auxiliary statements, respectively. Cramer distinguishes between input variables providing data to a program and local (body) variables used for computation. Any of them can be chosen as output variable at the end of the execution. The multiplication program used in Listing 4.1 can now be encoded as (0 0 1 3 5 8 1 3 2 1 4 3 4 5 9 9 2) which translates to

```

1  (0 0 1) ;;main statement → (:BLOCK AS0 AS1)
2  (3 5 8) ;;auxiliary statement 0 → (:ZERO V5)
3  (1 3 2) ;;auxiliary statement 1 → (:LOOP V3 AS2)
4  (1 4 3) ;;auxiliary statement 2 → (:LOOP V4 AS3)
5  (4 5 9) ;;auxiliary statement 3 → (:INC V5)

```

Listing 4.3: Another example for the JP Mapping

Cramer outlines some of the major problems of this representation, especially the strong positional epistasis¹⁶ – the strong relation of the meaning of an instruction to its position. This epistasis makes it very hard for the genetic operations to work efficiently, i.e., to prevent destruction of the genotypes passed to them.

¹⁶ We come back to positional epistasis in Section 4.8.1 on page 171.

The TB Mapping

The TB mapping is essentially the same as the JB mapping, but reduces these problems a bit. Instead of using the auxiliary statement method as done in JB, the expressions in the TB language are decoded recursively. The string $(0\ (3\ 5)(1\ 3\ (1\ 4\ (4\ 5)))\)$, for instance, expands to the program tree illustrated in Listing 4.3. Furthermore, Cramer restricts mutation to the statements near the fringe of the tree, more specifically, to leaf operators that do not require statements as arguments and to non-leaf operations with leaf statements as arguments. Similar restrictions apply to crossover.

4.4.2 Binary Genetic Programming

With their Binary Genetic Programming (BGP) approach [944], Keller and Banzhaf [945; 946; 947] further explore the utility of explicit genotype-phenotype mappings and neutral variations in the genotypes. They called the genes in their fixed-length binary string genome *codons* analogously to molecular biology where a codon is a triplet of nucleic acids in the DNA¹⁷, encoding one amino acid at most. Each codon corresponds to one symbol in the target language. The translation of the binary string genotype g into a string representing an expression in the target language works as follows:

1. $x \leftarrow \varepsilon$
2. Take the next gene (codon) g from g and translate it to the according symbol s .
3. If s is a valid continuation of x , set $x \leftarrow x \circ s$ and continue in step 2.
4. Otherwise, compute the set of symbols S that would be valid continuation of x .
5. From this set, extract the set of (valid) symbols S' which have the minimal Hamming distance¹⁸ to the codon g .
6. From S' take the symbol s' which has the minimal codon value and append it to x :
 $x \leftarrow x \circ s'$.

After this mapping, x can still be an invalid expression since there maybe were not enough genes in g so the phenotype is incomplete, for example $x = 3 * 4 - \sin(v*$. These incomplete sequences are fixed by consecutively appending symbols that lead to a quick end of an expression according to some heuristic.

The genotype-phenotype mapping of Binary Genetic Programming represents a $n : 1$ relation: Due to the fact that different codons may be replaced by the same approximation, multiple genotypes have the same phenotypic representation. This also means that there can be genetic variations induced by the mutation operation that do not influence the fitness. Such neutral variations are often considered as a driving force behind (molecular) evolution [948, 949, 950] and are discussed in Section 1.4.4 on page 46 in detail.

From the form of the genome we assume the number of corrections needed in the genotype-phenotype mapping (especially for larger grammars) will be high. This, in turn, could lead to very destructive mutation and crossover operations since if one codon is modified, the semantics of many subsequent codons may be influenced wildly. This issue is also discussed in Section 4.8.1 on page 170.

4.4.3 Gene Expression Programming

Gene Expression Programming (GEP) by Ferreira [951; 952; 953; 954; 955] introduces an interesting method for dealing with remaining unsatisfied function arguments at the end of the expression tree building process. Like BGP, Gene Expression Programming uses a genotype-phenotype mapping that translates fixed-length string chromosomes into tree phenotypes representing programs.

¹⁷ See Figure 1.14 on page 24 for more information on the DNA.

¹⁸ see Definition 28.6 on page 469

A gene in GEP is composed of a head and a tail [951] which are further divided into codons, where each codon directly encodes one expression. The codons in the head of a gene can represent arbitrary expressions whereas the codons in the tail can only stand for parameterless terms. This makes the tail a reservoir for unresolved arguments of the expressions in the head.

For each problem, the length h of the head is chosen as a fixed value, and the length of the tail t is defined according to Equation 4.2, where n is the arity (the number of arguments) of the function with the most arguments.

$$t = h(n - 1) + 1 \quad (4.2)$$

The reason for this formula is that we have h expressions in the head, each of them taking at most n parameters. An upper bound for the total number of arguments is thus $h * n$. From this number, $h - 1$ are already satisfied since all expressions in the head (except for the first one) themselves are arguments to expressions instantiated before. This leaves at most $h * n - (h - 1) = h * n - h + 1 = h(n - 1) + 1$ unsatisfied parameters. With this simple measure, incomplete expressions that require additional repair operations in BGP and most other approaches simply cannot occur.

For instance, consider the grammar for mathematical expressions with the terminal symbols $\Sigma = \{\sqrt{\cdot}, *, /, -, +, a, b\}$ given as example in [951]. It includes two variables, a and b , as well as five mathematical functions, $\sqrt{\cdot}$, $*$, $/$, $+$, and $-$. $\sqrt{\cdot}$ has the arity 1 since it takes one argument, the other four have arity 2. Hence, $n = 2$.




Fig. 4.14: A GPM example for Gene Expression Programming.

Figure 4.14 illustrates an example gene (with $h = 10$ and $t = h(2 - 1) + 1 = 11$) and its phenotypic representation of this mathematical expression grammar. A phenotype is built

by interpreting the gene as a level-order traversal¹⁹ of the nodes of the expression tree. In other words, the first codon of a gene encodes the root r of expression tree (here $+$). Then, all nodes in the first level (i. e., the children of r , here $\sqrt{\cdot}$ and $-$) are stored from left to right, then their children and so on. In the phenotypic representation, we have sketched the traversal order and numbered the levels. These level numbers are annotated to the gene but are neither part of the real phenotype nor the genotype. Furthermore, the division of the gene into head and tail is shown. In the head, the mathematical expressions as well as the variables may occur, while variables are the sole construction element of the tail.

In GEP, multiple genes form one genotype, thus encoding multiple expression trees. These trees may then be combined to one phenotype by predefined statements. It is easy to see that binary or integer strings can be used as genome, because the number of allowed symbols is known in advance.

This fixed mapping is also a disadvantage of Gene Expression Programming in comparison with the methods introduced later which have variable input grammars. On the other hand, there is the advantage that all genotypes can be translated to valid expression trees without requiring any corrections. Another benefit is that it seems to circumvent – at least partially – the problem of low causality from which the string-to-tree-GPM based approaches in often suffer. By modularizing the genotypes, potentially harmful influences of the reproduction operations are confined to single genes while others may stay intact. (See Section 4.8.1 on page 170 for more details.)

General Information

Areas Of Application

Some example areas of application of Gene Expression Programming are:

Application	References
Symbolic Regression and Function Synthesis	[956, 957, 958]
Data Mining and Data Analysis	[804, 805, 959, 960, 961, 962]
Electrical Engineering and Circuit Design	[963]
Machine Learning	[964, 965]
Geometry and Physics	[966, 967, 968]

Online Resources

Some general, online available ressources on Gene Expression Programming are:

http://www.gene-expression-programming.com/ [accessed 2007-08-19]	
Last update:	up-to-date
Description:	Gene Expression Programming Website. Includes publications, tutorials, and software.

4.4.4 Edge Encoding

Up until now, we only have considered how string genotypes can be transformed to more complex structures like trees. Obviously, genotype-phenotype mappings are not limited to this, but can work on tree genotypes as well. In [969], Luke and Spector present their *edge encoding approach* where the genotypes are trees (or forests) of expressions from a

¹⁹ http://en.wikipedia.org/wiki/Tree_traversal [accessed 2007-07-15]

graph-definition language. During the GPM, these trees are interpreted and construct the phenotypes, arbitrary directed graphs. Edge encoding is closely related to Gruau’s cellular encoding [900], which works on nodes instead of edges.

Each functions and terminals in edge encoding work on tuples (a, b) containing two node identifiers. Such a tuple represents a directed edge from node a to node b . The functions edit these tuples, add nodes or edges and thus, successively build the graph. Unlike normal Genetic Programming applications like symbolic regression, for instance, the nodes of trees in edge encoding are “executed” from top to bottom (pre-order) and pass control down to their children (from left to right). After an edge has been processed by a terminal node, it becomes permanent part of the graph constructed. In order to allow the construction of arbitrary graphs, an additional control structure, a stack of node identifiers, is used. Each node in the GP tree may copy this stack, modify this copy, and pass it to all of its children.

In their paper [969], Luke and Spector give multiple possible function and terminal sets for edge encoding. We provide a set that is sufficient to build arbitrary graphs in Table 4.1. Generally, each node receives an input edge tuple $E = (a, b)$ and a stack s which it can process. The two commands `labelE` and `labelN` in the table are no real functions but just here to demonstrate how nodes and edges can be enriched with labels and other sorts of information.

Operator	Children	Description
<code>double</code>	2	1. create a new edge $F = (a, b)$ 2. pass $E = (a, b)$ and the stack s to the first child 3. pass $F = (a, b)$ and the stack s to the second child
<code>bud</code>	2	1. create node c 2. create an edge $F = (b, c)$ 3. pass $E = (a, b)$ and the stack s to the first child 4. pass $F = (b, c)$ and the stack s to the second child
<code>split</code>	2	1. create a node c 2. change edge $E = (a, b)$ to $E = (a, c)$ 3. change edge $F = (c, b)$ 4. pass $E = (a, c)$ and the stack s to the first child 5. pass $F = (c, b)$ and the stack s to the second child
<code>loop</code>	2	1. create a new edge $F = (b, b)$ 2. pass $E = (a, b)$ and the stack s to the first child 3. pass $F = (b, b)$ and the stack s to the second child
<code>cut</code>	0	1. eliminate edge $E = (a, b)$
<code>nop</code>	0	1. make edge $E = (a, b)$ permanent
<code>push</code>	1	1. create a new node c 2. make a copy s' of the stack s 3. push c onto this copy s' 4. pass $E = (a, b)$ and the new stack s' to the child
<code>attach</code>	3	1. make a copy s' of the stack s 2. if s' is empty, create a new node and push it onto s' 3. pop the node c from the top of the stack s' 4. create two new edges $F = (a, c)$ and $G = (b, c)$ 5. pass $E = (a, b)$ and the new stack s' to the first child 6. pass $F = (a, c)$ and the new stack s' to the second child 7. pass $G = (b, c)$ and the new stack s' to the third child
<code>labelN</code>	1	1. label node b from edge $E = (a, b)$ with something 2. pass $E = (a, b)$ and the stack s to the child
<code>labelE</code>	1	1. label the edge $E = (a, b)$ with something 2. pass $E = (a, b)$ and the stack s to the child

Table 4.1: One possible operator set of edge encoding.




Fig. 4.15: An example for edge encoding.

In Figure 4.15, an example genotype for edge encoding is given. The nodes of this genotype are annotated with the parameters which are passed to them by their parent. The root receives an initial node tuple and an empty stack. Notice that we have replaced the node names a, b, c from Table 4.1 with running numbers and that new edges receive automatically a new name. At the bottom of the graphic, you find the result of the interpretation of the genotype by the GPM, a beautiful graph.

Edge encoding can easily be extended with automatically defined functions (as also shown in [969]) and gave the inspiration for Sinclair's node pair encoding method for evolving network topologies [872] (discussed in Section 23.1.1 on page 352). Vaguely related to such a graph generating approach are some of the methods for deriving electronic circuits by Lohn et al. [727] and Koza et al. [829] which you can find listed in the "Applications" tables in the general information sections.

4.5 Grammars in Genetic Programming

4.5.1 Introduction

We have learned that the most common genotypic and phenotypic representations in Genetic Programming are trees and also have discussed the reproduction operations available for tree-based genomes. In this discussion, we left one out important point: in many applications, reproduction cannot occur freely. Normally, there are certain restrictions to the structure and shape of the trees that must not be violated. Take our pet-example symbolic regression²⁰ for instance. If we have a node representing a division operation, it will take two arguments: the dividend and the divisor. One argument is not enough and a third argument is useless, as one can easily see in Figure 4.16.




Fig. 4.16: Example for valid and invalid trees in symbolic regression.

There are four general methods how to avoid invalid configurations under these limitations:

1. Compensate illegal configurations during the evaluation of the objective functions. This would mean, for example, that a division with no arguments could return 1, a division with only the single argument a could return a , and that superfluous arguments (like c in Figure 4.16) would simply be ignored.
2. A subsequent repair algorithm could correct errors in the tree structure that have been introduced during reproduction.
3. Using additional checking and refined node selection algorithms, we can ensure that only valid trees are created during the reproduction cycle.
4. With special genotype-phenotype mappings, we can prevent the creation of invalid trees from the start.

In this section, we will introduce some general methods of enforcing valid configurations in the phenotypes, mostly regarding the fourth approach. A very natural way to express structural and semantic restrictions of a search space are formal grammars which are elaborated on in Section 29.3 on page 501. Genetic Programming approaches that limit their phenotypes (the trees) to sentences of a formal language are subsumed under the topic of Grammar-guided Genetic Programming (GGGP).

4.5.2 Trivial Approach

Standard Genetic Programming as introduced by Koza [65] already inherently utilizes simple mechanisms to ensure the correctness of the tree structures. These mechanisms are rather trivial, though, and should not be counted to the family of GGGP approaches, but are mentioned here for the sake of completeness.

In Standard Genetic Programming, all expressions have exactly the same type. Applied to symbolic regression, this means that, for instance, all constructs will be real-valued or

²⁰ See Section 22.1 on page 341.

```

1 <Z> ::= (<Z> + <Z>)
2 <Z> ::= (<Z> - <Z>)
3 <Z> ::= (<Z> * <Z>)
4 <Z> ::= (<Z> / <Z>)
5 <Z> ::= (sin <Z>)
6 <Z> ::= X

```

Listing 4.4: A trivial symbolic regression grammar.

return real values. If logical functions like multiplexers are grown, all entities will be Boolean-valued, and so on. For each possible tree node type, we just need to specify the exact amount of children. This approach corresponds to a context-free grammar²¹ with a single non-terminal symbol which is expanded by multiple rules. Listing 4.4 illustrates such a trivial grammar $G = (N, \Sigma, P, S)$ in Backus-Naur Form (BNF)²². Here, the non-terminal symbol is Z ($N = \{Z\}$), the terminal symbols are $\Sigma = \{(\,,\,), +, -, *, /, \text{sin}, X\}$, and six different productions are defined. The start symbol is $S = Z$.

Standard Genetic Programming does not utilize such grammars directly. Rather, they are hard-coded in the reproduction operators or are represented in fixed internal data structures.

Here we should mention that illegal configurations can also rise at runtime from semantics. In symbolic regression, a division operation is invalid if the divisor is zero, for instance. The same goes for logarithms, or a tangent of $(n + \frac{1}{2})\pi \forall n \in \mathbb{Z}$. All four approaches for enforcing a proper tree structure previously introduced cannot prevent such errors from the start. Therefore, the function set (the possible inner nodes of the trees) need to ensure the property of *closure* as defined by Koza [65].

Definition 4.1 (Closure). If a function set N has the property *closure*, it ensures that all possible values are accepted as parameter by any function.

Closure is especially important in approaches like symbolic regression, and can easily be achieved by redefining the mathematical functions for special cases, like setting $\frac{a}{0} = a \forall a \in \mathbb{R}$, for instance. It does, however, not consider the tree structure itself – the number of arguments still needs to be sufficient.

4.5.3 Strongly Typed Genetic Programming

The strongly typed Genetic Programming (STGP) approach developed by Montana [970; 971; 972] is still very close to Standard Genetic Programming. With strongly typed Genetic Programming, it becomes possible to use typed data structures and expressions in Genetic Programming. Hence, the issue of well-typedness arises, as illustrated in Figure 4.17.




Fig. 4.17: Example for valid and invalid trees in typed Genetic Programming.

²¹ see Section 29.3.2 on page 503 for details

²² The Backus-Naur form is discussed in Section 29.3.4 on page 504.

As already mentioned in Section 4.5.2 on page 147, in Standard Genetic Programming such errors are circumvented by only using representations that are type-safe per definition. In symbolic regression, for instance, only functions and variables which are real-typed are allowed, and in the evolution of logic functions only Boolean-valued expressions will be admitted. Thus, inconsistencies like in Figure 4.17 are impossible.

In STGP, a tree genome is used which permits different data types that are not assignment-compatible. One should not mistake STGP for a fully grammar-guided approach, since it uses rules still based on an implicit, hard-coded internal grammar which are built in the bootstrap phase of the GP system. However, it represents clearly a method to shape the individuals according to some validity constraints.

These constraints are realized by modified reproduction operations that use *types possibilities tables* which denote which types for expressions are allowed in which level of a tree (individual). The creation and mutation operators now return valid individuals per default. Recombination still selects the node to be replaced in the first parent randomly, but the sub-tree in the second parent which should replace this node is selected in a way that ensures that the types match. If this is not possible recombination either returns the parents or an empty set.

STGP also introduces interesting new concepts like generic functions and data types, very much like in Ada or C [972] and hierarchical type systems, comparable to object-oriented programming in their inheritance structure [973]. This way, STGP increases the reusability and modularity in GP which is needed for solving more complex problems [883, 925].

4.5.4 Early Research in GGGP

Research steps into grammatically driven program evolution can be traced to the early 1990s where Antonisse [974] developed his Grammar-based Genetic Algorithm. As genome, he used character strings representing sentences in a formal language defined by a context-free grammar. Whenever crossover was to be performed, these strings were parsed into the derivation trees²³ of that grammar. Then, recombination was applied in the same way as in tree-based systems. This parsing was the drawback of the approach, leading to two major problems: First, it slows down the whole evolution since it is an expensive operation. Secondly, if the grammar is ambiguous, there may be more than one derivation tree for the same sentence [810]. Antonisse's early example was succeeded by other researchers like Stefanski [975], Roston [976], and Mizoguchi et al. [977].

In the mid-1990s [810, 978], more scientists began to concentrate on this topic. The LOGENPRO system developed by Wong and Leung [979], [980, 981, 982, 983, 984, 825] used PROLOG Definite Clause Grammars to derive first-order logic programs. A GP system proposed by Whigham [985; 986; 987; 988; 989] applied context-free grammars in order to generate populations of derivation trees. This method additionally had the advantage that it allowed the user to bias the evolution into the direction of certain parts of the grammar [988]. Geyer-Schulz [932] derived a similar approach, differing mainly in the initialization procedure [990, 810], for learning rules for expert systems. The Genetic Programming Kernel (GPK) by Hörner [991] used tree-genomes where each genotype was a deviation tree generated from a BNF definition.

4.5.5 Gads 1

The Genetic Algorithm for Deriving Software 1 (Gads 1) by Paterson and Livesey [992; 993] was one of the basic research projects that paved the way for other, more sophisticated, approaches like Grammatical Evolution. Like the Binary Genetic Programming system by Keller and Banzhaf [945], it uses a clear distinction between genotype and phenotype. The genotypes $g \in \mathbb{G}$ in Gads are fixed-length integer strings which are transformed to character

²³ An elaboration on derivation trees can be found in Section 29.3.3 on page 503.

```

1  (0) <expr>      ::= <expr> <op> <expr>
2  (1) <expr>      ::= (<expr> <op> <expr>)
3  (2) <expr>      ::= <pre-op> (<expr>)
4  (3) <expr>      ::= <var>
5
6  (4) <op>         ::= +
7  (5) <op>         ::= -
8  (6) <op>         ::= /
9  (7) <op>         ::= *
10
11 (8) <pre-op>      ::= log
12 (9) <pre-op>      ::= tan
13 (10) <pre-op>     ::= sin
14 (11) <pre-op>     ::= cos
15
16 (12) <var>        ::= X
17
18 (13) <func>       ::= double func(double x){
19                             return <expr>;
20                             }

```

Listing 4.5: A simple grammar for C functions that could be used in Gads.

string phenotypes $prspel \in \mathbb{X}$ (representing program syntax trees) by a genotype-phenotype mapping (see Section 3.6 on page 122). Because of this genome, Gads can use a conventional genetic algorithm engine²⁴ to evolve the solution candidates.

Gads receives a context-free grammar $G = (N, \Sigma, P, S)$ specified in Backus-Naur form as input. In Binary Genetic Programming, the genome encodes the sequence of terminal symbols of the grammar directly. Here, a genotype specifies the sequence of the productions to be applied to build a sentence of terminal symbols.

Although Gads was primarily tested with LISP S-expressions, it can evolve sentences on basis of all possible BNF grammars. For the sake of coherence with later sections, we use a grammar for simple mathematical functions in C as example. Here, the set of possible terminals is $\Sigma = \{\text{sin}, \text{cos}, \text{tan}, \text{log}, +, -, *, /, \text{X}, ()\}$ and as non-terminal symbols we use $N = \{\text{expr}, \text{op}, \text{pre-op}, \text{func}\}$. The starting symbol is $S = \text{func}$ and the set of productions P is illustrated in Listing 4.5.

In the BNF grammar definitions for Gads, the “|” symbol commonly denoting alternatives is not used. Instead, multiple productions may be defined for the same non-terminal symbol.

Every gene in a Gads genotype contains the index of the production in G to be applied next. For now, let us investigate the genotype $g = (2, 0, 12, 5, 5, 13, 10)$ as example. If the predefined start symbol is **func**, we would start with the phenotype string x_1

```

1 double func(double x){
2     return <expr>;
3 }

```

The first gene in g , 2, leads to the application of rule (2) to x_1 and we obtain x_2 :

```

1 double func(double x){
2     return <pre-op> (<expr>);
3 }

```

The next gene is 0, which means that we will use production (0)). There is a **expr**-non-terminal symbol in x_2 , so we get x_3 as follows:

²⁴ Gads 1 uses the genetic algorithm C++ class library GAGS (<http://geneura.ugr.es/GAGS/> [accessed 2007-07-09]) release 0.95.

```

1 double func(double x){
2     return <pre-op> (<expr> <op> <expr>);
3 }

```

Now comes the next gene with allele 12²⁵. We cannot apply rule (12) since no `var`-symbol can be found in x_3 – we simply ignore this gene and set $x_3 = x_4$. The following gene with value 5 translates the symbol `op` to `-` and we obtain for x_5 :

```

1 double func(double x){
2     return <pre-op> (<expr> - <expr>);
3 }

```

The next two genes, 5 and 13, must again be ignored ($x_7 = x_6 = x_5$). Finally, the last gene with the allele 10 resolves the non-terminal `pre-op` and we get for x_8 :

```

1 double func(double x){
2     return sin (<expr> - <expr>);
3 }

```

For the remaining two `expr` non-terminal symbols no rule is defined in the genotype g . There are several ways for dealing with such incomplete resolutions. One would be to exclude the individual from evaluation/simulation and to give it the lowest possible objective values directly. Gads instead uses simple default expansion rules. In this example, we could translate all remaining `exprs` to `vars` and these subsequently to `x`. This way we obtain the resulting function below.

```

1 double func(double x){
2     return sin (X - X);
3 }

```

One of the problems in Gads is the unacceptable large number of introns²⁶ [44] caused by the encoding scheme. Many genes will not contribute to the structure of the phenotype since they encode productions that cannot be executed (like allele 12 in the example genotype g) because there are no matching non-terminal symbols. This is especially the case in “real-world” applications where the set of non-terminal symbols N becomes larger.

With the Gads system, Paterson paved the way for many of the advanced techniques described in the following sections.

4.5.6 Grammatical Evolution

Like Gads, Grammatical Evolution²⁷ (GE), developed by Ryan et al. [978], creates expressions in a given language by iteratively applying the rules of a grammar specified in the Backus-Naur form [978, 994, 995].

In order to discuss how Grammatical Evolution works, we re-use the example of C-style mathematical functions [978] from Section 4.5.5. Listing 4.6 specifies the according rules using a format which is more suitable for grammatical evolution.

There are five rules in the set of productions P , labeled from **A** to **E**. Some of the rules have different options (separated by `|`). In each rule, options are numbered started with 0. When the symbol `<exp>` for example is expanded, for example, there are four possible results (0-3). The shape of the sentences produced by the grammar depends on these choices.

Like in Gads, the genotypes in GE are numerical strings. These strings encode the indices of the options instead of the productions themselves. In Gads, each option was treated as a single production because of the absence of the “`|`” operator. The idea of Grammatical Evolution is that it is already determined which rules must be used by the non-terminal

²⁵ An allele is a value of specific gene, see Definition 1.24 on page 24.

²⁶ Introns are genes or sequences of genes (in the genotype) that do not contribute to the phenotype or its behavior, see Definition 3.2 on page 118

²⁷ http://en.wikipedia.org/wiki/Grammatical_evolution [accessed 2007-07-05]

```

1 (A) <expr> ::= <expr> <op> <expr> (0)
2           | (<expr> <op> <expr>) (1)
3           | <pre-op> (<expr>) (2)
4           | <var> (3)
5
6 (B) <op> ::= + (0)
7         | - (1)
8         | / (2)
9         | * (3)
10
11 (C) <pre-op> ::= log (0)
12           | tan (1)
13           | sin (2)
14           | cos (3)
15
16 (D) <var> ::= X (0)
17
18 (E) <func> ::= double func(double x){
19           return <expr>;
20           } (0)

```

Listing 4.6: A simple grammar for C functions that could be used by GE.

symbol to be expanded and we only need to decide which option of this rule is to be applied. Therefore, the number of introns is dramatically reduced compared to Gads.

The variable-length string genotypes of Grammatical Evolution can again be evolved using genetic algorithms [978, 996] (like in Gads) or with other techniques, like Particle Swarm Optimization [997, 998] or Differential Evolution [999]. As illustrated in Figure 4.18, a Grammatical Evolution system consists of three components: the problem definition (including the means of evaluating a solution candidate), the grammar that defines the possible shapes of the individuals, and the search algorithm that creates the individuals [1000].




Fig. 4.18: The structure of a Grammatical Evolution system [1000].

An Example Individual

We get back to our mathematical C function example grammar in Listing 4.6. As already said, a genotype $g \in \mathbb{G}$ is a variable-length string of numbers that denote the choices to be taken whenever a non-terminal symbol from N is to be expanded and more than one option is available (as in the productions (A), (B), and (C)). The start symbol, $S = \text{func}$ does not need to be encoded since it is predefined. Rules with only one option do not consume information from the genotype. The processing of non-terminal symbols uses a depth-first order [978], so resolving a non-terminal symbol ultimately to terminal symbols has precedence before applying an expansion to a sibling.

Let us assume we have settled for bytes as genes in the genome. As we may have less than 256 options, we apply modulo arithmetic to get the index of the option. This way, the sequence $g = (193, 47, 51, 6, 251, 88, 63)$ is a valid genotype. According to our grammar, the first symbol to expand is $S = \text{func}$ (rule (E)) where only one option is available. Therefore, all phenotypes will start out like

```
1 double func(double x){
2   return <expr>;
3 }
```

The next production we have to check is (A), since it expands **expr**. This production has four options, so taking the first number from the genotype g , we get $193 \bmod 4 = 1$ which means that we use option (1) and obtain

```
1 double func(double x){
2   return (<expr> <op> <expr>);
3 }
```

As **expr** appears again, we have to evaluate rule (A) once more. The next number, 47, gives us $47 \bmod 4 = 3$ so option (3) is used.

```
1 double func(double x){
2   return (<var> <op> <expr>);
3 }
```

var is expanded by rule (D) where only one result is possible:

```
1 double func(double x){
2   return (X <op> <expr>);
3 }
```

Subsequently, **op** will be evaluated to ***** since $51 \bmod 4 = 3$ (rule (B) (3)) and **expr** becomes **pre-op(<expr>)** because $6 \bmod 4 = 2$ (production (A) (2)). Rule (C) (3) then turns **pre-op** into **cos** since $251 \bmod 4 = 3$. **expr** is expanded to **<expr> <op> <expr>** by (A) (0) because $88 \bmod 4 = 0$. The last gene in our genotype is 63, and thus rule (A) (3) ($63 \bmod 4 = 3$) transforms **expr** to **<var>** which then becomes **X**.

```
1 double func(double x){
2   return (X * cos(X <op> <expr>));
3 }
```

By now, the numbers available in g are exhausted and we still have non-terminal symbols left in the program. As already outlined earlier, there are multiple possible approaches how to proceed in such a situation:

1. Mark g as invalid and give it a reasonably bad fitness.
2. Expand the remaining non-terminals using default rules (i. e., we could say the default value for **expr** is **X** and **op** becomes **+**),
3. or wrap around and restart taking numbers from the beginning of g .

The latter method is applied in Grammatical Evolution. It has the disadvantage that it can possibly result in an endless loop in the genotype-phenotype translation, so there should be a reasonable maximum for the iteration steps after which we fall back to default rules.

In the example, we will proceed by expanding **op** according to (B) (1) since $193 \bmod 4 = 1$ and obtain **-** (minus). The next gene gives us $47 \bmod 4 = 3$ so the last **expr** will become a **<var>** and finally our phenotype is:

```
1 double func(double x){
2   return (X * cos(X - X));
3 }
```

Note that if the last gene 63 was missing in g , the “restart method” which we have just described would produce an infinite loop, because the first non-terminal to be evaluated whenever we restart taking numbers from the front of the genome then will always be **expr**. In this example, we are lucky and this is not the case since after wrapping at the genotype end, a **pre-op** is to be resolved. The gene 193 thus is an index into rule **A** at its first usage and an index into production **C** in the second application.

Initialization

Grammatical Evolution uses an approach for initialization similar to ramped half-and-half²⁸, but on basis of derivation trees²⁹. Therefore, the numbers of the choices made during a random grammatical rule expansion beginning at the start symbol are recorded. Then, a genotype is built by reversing the modulo operation, i. e., finding a number that produces the same number as recorded when modulo-divided for each gene. The number of clones is subsequently reduced and, optionally, the single-point individuals are deleted.

General Information

Areas Of Application

Some example areas of application of Grammatical Evolution are:

Application	References
Mathematical Problems (vs. Standard Genetic Programming: [65])	[996, 978]
Automated Programming	[995, 1001, 1002]
Robotics (vs. Standard Genetic Programming: [849, 853])	[856, 1003]
Economics and Finance (vs. Standard Genetic Programming: [837, 838])	[839, 1004, 1005, 1006]

There even exists an approach called “Grammatical Evolution by Grammatical Evolution” ($(GE)^2$, [1007]) where the grammar defining the structure of the solution candidates itself is co-evolved with the individuals.

Conferences, Workshops, etc.

Some conferences, workshops and such and such on Grammatical Evolution are:

GEWS: Grammatical Evolution Workshop
http://www.grammatical-evolution.com/gews.html [accessed 2007-09-10]
History: 2004: Seattle, WA, USA, see [1008]
2003: Chicago, IL, USA, see [1009]
2002: New York, NY, USA, see [566]

Online Resources

Some general, online available ressources on Grammatical Evolution are:

http://www.grammatical-evolution.com/ [accessed 2007-07-05]

²⁸ An initialization method of standard, tree-based Genetic Programming that creates a good mixture of various tree shapes [65], see Section 4.3.1 on page 133 for more details.

²⁹ see Section 29.3.3 on page 503

Last update: up-to-date

Description: Grammatical Evolution Website. Includes publications, links, and software.

4.5.7 Gads 2

In Gads 2, Paterson [44] uses the experiences from Gads 1 and the methods of the Grammatical Evolution approach to tackle context-sensitive grammars with Genetic Programming. While context-free grammars are sufficient to describe the syntax of a programming language, they are not powerful enough to determine if a given source code is valid. Take for example the C snippet

```
1 char i;
2 i = 0.5;
3 y = 1;
```

It is obviously not a well-typed program although syntactically correct. Context-sensitive grammars³⁰ allow productions like $\alpha A \beta \rightarrow \alpha \gamma \beta$ where $A \in N$ is a non-terminal symbol, and $\alpha, \beta, \gamma \in V^*$ are concatenations of arbitrary many terminal and non-terminal symbols (with the exception that $\gamma \neq \varepsilon$, i. e., it must not be the empty string). Hence, it is possible to specify that a value assignment to a variable must be of the same type as the variable and that the variable must have previously been declared with a context-sensitive grammar. Paterson argues that the application of existing approaches like two-level grammars and standard attribute grammars³¹ in Genetic Programming is infeasible [44] and introduces an approach based on *reflective attribute grammars*.

Definition 4.2 (Reflective Attribute Grammar). A reflective attribute grammar (rag³²) [44] is a special form of attribute grammars. When expanding a non-terminal symbol with a rag production, the grammar itself is treated as an (inherited) attribute. During the expansion, it can be modified and is finally passed on to the next production step involving the newly created nodes.

The transformation of a genotype $g \in \mathbb{G}$ into a phenotype using a reflective attribute grammar r resembles Grammatical Evolution to some degree. Here we discuss it with the example of the recursive expansion of the symbol s :

1. Write the symbol s to the output.
2. If $s \in \Sigma$, i. e., s is a terminal symbol, nothing else is to do – return.
3. Use the next gene in the genotype g to choose one of the alternative productions that have s on their left hand side. If g is exhausted, choose the default rule.
4. Create the list of the child symbols $s_1 \dots s_n$ according to the right-hand side of the production.
5. For $i = 1$ to n do
 - a) Resolve the symbol i , passing in s_i , r , and g .
 - b) If needed, modify the grammar r according to the semantics of s and s_i .

Item 5 is the main difference between Gads 2 and Grammatical Evolution. What happens here depends on the semantics in the rag. For example, if a non-terminal symbol that declares a variable x is encountered, a new terminal symbol κ is added to the alphabet Σ that corresponds to the name of x . Additionally, the rule which expands the non-terminal

³⁰ See Section 29.3.2 on page 503 where we discuss the Chomsky Hierarchy of grammars.

³¹ See Section 29.3.6 on page 505 for a discussion of attribute grammars.

³² Notice that the shortcut of this definition *rag* slightly collides with the one of Recursive Adaptive Grammars (*RAG*) introduced by Shutt [1010] and discussed in Section 29.3.8 on page 508, although their letter cases differ. To the knowledge of the author, rags are exclusively used in Gads 2.

symbol that stands for variables of the same type now is extended by a new option that returns κ . Thus, the new variable becomes available in the subsequent code.

Another difference to Grammatical Evolution is the way the genes are used to select an option in item 3. GE simply uses the modulo operation to make its choice. Assume we have genotypes where each gene is a single byte and encounter a production with four options while the next gene has the value 45. In Grammatical Evolution, this means to select the second option since $45 \bmod 4 = 1$ and we number the alternatives beginning with zero. Gads 2, on the other hand, will divide the range of possible alleles into four disjoint intervals of (approximately) equal size $[0, 63]$, $[64, 127]$, $[128, 191]$, $[192, 255]$ where 45 falls clearly into the first one. Thus, Gads 2 will expand the first rule.

The advantage of Gads 2 is that it allows to grow valid sentences according to context-sensitive grammars. It becomes not only possible to generate syntactically correct but also well-typed source code for most conventional programming languages. Its major drawback is that it has not been realized fully. The additional semantics of the production expansion rule 5b have not been specified in the grammar or in an additional language as input for the Genetic Programming system but are only exemplarily realized in a hard-coded manner for the programming language S-Algol [1011]. The experimental results in [44], although successful, do not provide substantial benefits compared to the simpler Grammatical Evolution approach.

Gads 2 shows properties that we also experienced in the past: Even if constructs like loops, procedure calls, or indexed access to memory are available, the chance that they are actually used in the way in which we would like them to be used is slim. Genetic Programming of real algorithms in a high-level programming language-like syntax exhibits a high affinity to employ rather simple instructions while neglecting more powerful constructs. Good fitness values are often reached with overfitting only.

Like Grammatical Evolution, the Gads 2 idea can be realized with arbitrary genetic algorithm engines. Paterson [44] uses the Java-based evolutionary computation system ECJ by Luke et al. [1012] as genetic algorithm engine in his experiments.

4.5.8 Christiansen Grammar Evolution

Christiansen Grammar, which you can find described in Section 29.3.9 on page 509, have many similarities to the reflective attribute grammars used in Gads 2. They are both Extended Attribute Grammars³³ and the first attribute of both grammars is an inherited instance of themselves. Christiansen Grammars are formalized and backed by comprehensive research since being developed back in 1985 by Christiansen [1013].

Building on their previous work de la Cruz Echeandía et al. [1014] place the idea of Gads 2 on the solid foundation of Christiansen Grammars with their Christiansen Grammar Evolution approach (CGE) [1015]. They tested their system for finding logic function identities with constraints on the elementary functions to be used. Instead of elaborating on this experiment, let us stick with the example of mathematical functions in C for the sake of simplicity.

In Listing 4.7 we define the productions P of a Christiansen Grammar that extends the examples from before by the ability of creating and using local variables. Three new rules (F), (G), and (H) are added, and the existing ones have been extended with attributes.

The non-terminal symbol `expr` now receives the inherited attribute $\mathbf{g}$ which is the (Christiansen) grammar to be used for its expansion. The $\downarrow$ (arrow down) indicates inherited attribute values that are passed down from the parent symbol, whereas $\uparrow \mathbf{a}$ (arrow up) identifies an attribute value $\mathbf{a}$ synthesized during the expansion of a symbol and passed back to the parent symbol.

The start symbol S is still `func`, but the corresponding production (E) has been complemented by a reference to the new non-terminal symbol `stmt` (line 19). The symbol `stmt`

³³ See Section 29.3.7 on page 507 for more information on such grammars.

```

1  (A) <expr ↓g>      ::= <expr↓g> <op↓g> <expr↓g>      (0)
2                      | (<expr ↓g> <op ↓g> <expr ↓g>) (1)
3                      | <pre-op ↓g> (<expr ↓g>)         (2)
4                      | <var ↓g>                        (3)
5
6  (B) <op ↓g>        ::= "+"                            (0)
7                      | "-"                            (1)
8                      | "/"                            (2)
9                      | "*"                            (3)
10
11 (C) <pre-op ↓g>     ::= "log"                          (0)
12                      | "tan"                          (1)
13                      | "sin"                          (2)
14                      | "cos"                          (3)
15
16 (D) <var ↓g>        ::= "X"                            (0)
17
18 (E) <func ↓g1>      ::= "double_func(double_x){"
19                      <stmt ↓g1 ↑g2>
20                      "return_<expr ↓g2> "; "
21                      }                                (0)
22
23 (F) <stmt ↓g0 ↑g2>  ::= <stmt ↓g0 ↑g1><stmt ↓g1 ↑g2> (0)
24                      | <new-var ↓g0 ↑g2>              (1)
25                      | <assign ↓g0 ↑g2>                (2)
26
27 (G) <new-var ↓g ↑g+new-rule> ::=
28     "double_<alpha-list ↓g ↑w> "=0;"                (0)
29     where <new-rule> is <var ↓g> ::= w
30
31 (H) <assign ↓g ↑g> ::= <var ↓g> "=" <expr ↓g> ";"    (0)

```

Listing 4.7: A Christiansen grammar for C functions that use variables.

has two attributes: an inherited (input) grammar g_0 and a synthesized (output) grammar g_2 . We need to keep that in mind when discussing the options possible for its resolution. A `stmt` symbol can either be expanded to two new `stmts` in option (0), a variable declaration represented by the non-terminal symbol `new-var` as option (1), or to a variable assignment (symbol `assign`) in option (2). Most interesting here is option (1), the variable declaration.

The production for `new-var`, labeled (G), receives the grammar g as input. The synthesized attribute it generates as output is g extended by a new rule `new-rule`. The name of the new variable is a string over the Latin alphabet. In order to create this string, we make use of the non-terminal symbol `alpha-list` defined in Listing 29.12 on page 509. `alpha-list` inherits a grammar as first attribute, generates a character string w , and also synthesizes it as output. Production (G) uses this value w in order to build its output grammar. It creates a new rule (see line 29) which extends the production (D) by a new option. `var` can now be resolved to either `X` or to one of the new variables in subsequent expansions of `expr` because the synthesized grammar is passed up to `stmt` and from there to all subsequent statements (see rule (F) option (0)) and even by the returned expression in line 20. It should be mentioned that this example grammar does not prevent name collisions of the identifiers, since `X`, for instance, is also a valid expansion of `new-var`.

With this grammar, a Christiansen Grammar Evolution system would proceed exactly as done in Section 4.5.6 on page 151.

4.5.9 Tree-Adjoining Grammar-guided Genetic Programming

A different approach to Grammar-guided Genetic Programming has been developed by Nguyen [260] with his Tree-Adjoining Grammar-guided Genetic Programming (TAG3P) system [1016, 1017, 1018, 1019, 1020]. Instead of using grammars in the Backus-Naur Form or one of its extensions as done in the aforementioned methods, it bases on tree-adjoining grammars (TAGs) which are introduced in Section 29.3.10 on page 510.

An Example TAG grammar

A tree-adjoining grammar can be defined as quintuple $G = (N, \Sigma, A, I, S)$ where N are the non-terminal, Σ contains the terminal symbols, and S is the start symbol. TAGs support two basic operations: adjunction and substitution. For these operations, blueprint trees are provided in the set of auxiliary and initial trees respectively (A and I). Substitution is quite similar to expansion in BNF, the root of an initial tree replaces a leaf with the same label in another tree. A tree β to be used for adjunction has at least one leaf node ν (usually marked with an asterisk $*$) with the same label as its root. It is injected into another tree by replacing a node with (again) that label whose children are then attached to ν .




Fig. 4.19: An TAG realization of the C-grammar of Listing 4.6.

Let us take a look back on the tree-adjoining representation of our earlier example grammar G in Listing 4.6 on page 152 for mathematical functions in C. Figure 4.19 illustrates one possible realization of G as TAG. The productions are divided into the set of initial trees I , which are used in substitution operations, and the auxiliary trees A needed by the adjunction operator. Again, the start symbol is `func` – this time however it identifies a tree in I . We additionally have annotated the trees with the index of the corresponding rule in Listing 4.6. It is possible that we need to build multiple TAG trees for one BNF rule, as done with rule 1 which is reflected in the two auxiliary trees β_1 and β_2 . The rules 3 and 12 on the other hand have been united into one initial tree for the purpose of simplicity (It could have been done in the BNF in the same way).

Like the other grammar-guided methods, the TAG3P approach uses a genotype-phenotype mapping. The phenotypes are, of course, trees that comply with the input tree-adjoining grammar. The genotypes being evolved are derivation trees that work on this grammar too. Derivation trees illustrate the way the productions of a grammar are applied in order to derive a certain sentence, as discussed in Section 29.3.3 on page 503.

Derivation Trees

For tree-adjoining grammars, there exist different types of derivation trees [260]. In the method of Weir [1021], they are characterized as object trees where the root is labeled with an S -type initial tree (i. e., the start symbol) and all other trees are labeled with the names of auxiliary trees. Each connection from a parent p to a child node c is labeled with the index of the node in p being the center of the operation. Indices are determined by numbering the non-terminal nodes according to a preorder traversal³⁴. The number of adjunctions performed with each node is limited to one. Substitution operations are not possible with Weir’s approach. Joshi and Schabes [1022] introduce an extension mitigating this problem. In their notation (not illustrated here) a solid connection between two nodes in the derivation tree stands for adjunction, whereas a broken line denotes a substitution.

In TAG3P, Nguyen [260] uses a restricted form of such TAG derivation trees where adjunction is not permitted to (initial) trees used for substitution. This essentially means that all adjunctions are performed before any substitutions. With this definition, substitutions become basically in-node operations. We simply attach the nodes substituted into a tree as list of lexemes (here terminal symbols) to the according node of a derivation tree.

Example Mapping: Derivations Tree $\rightarrow$ Tree

Figure 4.20 outlines some example mappings from derivation trees on the left side to sentences of the target languages (displayed as trees) on the right side. In Figure 4.19, we have annotated some of the elementary trees with α or β and numbers which we will use here. The derivation tree α_1 , for example, represents the initial production for the starting symbol. In addition, we have attached the preorder index to each node of the trees α_1 , β_3 , and β_5 . In the next tree we show how the terminal symbols `x` and `+` can be substituted into β_3 . In the corresponding derivation tree, they are simply attached as a list of lexemes. A similar substitution can be performed with β_5 , where `sin` is attached as terminal symbol.

In the fourth example, the second derivation tree is adjoined to the first one. Since it replaces the node with the preorder index 1, the connection from β_3 to α_1 is labeled with 1. Finally, in the fifth example, the third derivation tree is adjoined. We use the rule for `preops` to replace the node number 3 (according to preorder) in the second derivation in its adjoined state.

As you can see, all initial trees as well as all trees derived from them are always valid sentences of the grammar. This means that we can remove any of the derivation steps and still get valid phenotypes. Thus, we can evaluate the share of the fitness clubbed by every single modification by evaluating the resulting phenotypes with and without it.

³⁴ http://en.wikipedia.org/wiki/Tree_traversal [accessed 2007-07-18]




Fig. 4.20: One example genotype-phenotype mapping in TAG3P.

Summary

Tree-Adjoining Grammar-guided Genetic Programming is a different approach to Grammar-guided Genetic Programming which has some advantages compared with the other methods. One of them is the increased domain of locality. All nodes of a derivation tree stay accessible for the reproduction operations. This becomes interesting when modifying nodes “without side effects to other regions of the resulting trees”. If we, for example, toggle one bit in a Grammatical Evolution-based genotype, chances are that the meaning of all subsequent genes change and the tree resulting from the genotype-phenotype mapping will be totally different from its parent. In TAG3P, this is not the case. All operations can, at most, influence the node they are applied to and its children. Here, the principle of strong causality holds since small changes in the genotype lead to small changes in the phenotype. On the other hand, some of these positive effects may also be reached more easily with the wrapping and lifting operations for Genetic Programming introduced in this book in Section 4.3.7 on page 136 and Section 4.3.8. The reproduction operations of TAG3P become a little bit more complicated. When performing crossover, for instance, we can only exchange *compatible* nodes. We cannot adjoin the tree α_1 in Figure 4.20 with itself, for example.

General Information

Areas Of Application

Some example areas of application of Tree-Adjoining Grammar-guided Genetic Programming are:

Application	References
Symbolic Regression and Function Synthesis	[1017, 1018, 1019]
Mathematical Problems	[1023, 1024]

Online Resources

Some general, online available resources on Tree-Adjoining Grammar-guided Genetic Programming are:

http://sc.snu.ac.kr/SCLAB/Research/publications.html [accessed 2007-09-10]	
Last update:	up-to-date
Description:	Publications of the Structural Complexity Laboratory of the Seoul National University, includes Nguyen’s papers about TAG3P

4.6 Linear Genetic Programming

In the beginning of this chapter, we have learned that the major goal of Genetic Programming is to find programs that solve a given set of problems. We have seen that tree genomes are suitable to encode such programs and how the genetic operators can be applied to them.

Trees are however not the only way of representing programs. Matter of fact, a computer processes them in form of a sequence of instructions instead. This sequence may contain branches in form of jumps to other places in the program. Every possible flowchart describing the behavior of a program can be translated into such a sequence. It is therefore only natural that the first approach to automated program generation by Friedberg [783] in the 1950s used a fixed-length instruction sequence genome [783, 784].

The area of Genetic Programming that works with such instruction string genomes is called *linear Genetic Programming* (LGP) [1025, 1026, 1027, 99, 1028, 1029, 931] in contrast to the traditional *tree-based Genetic Programming* (TGP). It can furthermore be distinguished from approaches like Grammatical Evolution (see Section 4.5.6 on page 151) by

the fact that strings there are just genotypic, intermediate representations that encode the program trees. In LGP, they are the center of the whole evolution and contain the program code directly.

The advantage of linear Genetic Programming lies in the easy evaluation of the evolved algorithms. Its structure furthermore eases limiting the runtime in individual evaluation and even to simulate parallelism. Their drawback is that we cannot simply reuse the genetic operators for variable-length string genomes (discussed in Section 3.5 on page 121) which randomly insert, delete, or toggle bits, is thus not really feasible [1029].

We must visualize, for example, that the alternatives and loops which we know from high-level programming languages are mapped to conditional and unconditional jump instructions in machine code. These jumps target to either absolute or relative addresses inside the program. Let us consider the insertion of a single, new command into the instruction string, maybe as result of a mutation or recombination operation. If we do not perform any further corrections after this insertion, it is well possible that the resulting shift of the absolute addresses of the subsequent instructions in the program invalidates the control flow and renders the whole program useless. This issue is illustrated in Fig. 4.21.1. Nordin et al.




Fig. 4.21.1: Inserting into an instruction string.




Fig. 4.21.2: Inserting in a tree representation.

Fig. 4.21: The impact of insertion operations in Genetic Programming

[1030; 1030] point out that standard crossover is highly disruptive. Even though the subtree crossover in tree-genomes is shown to be [935] not very efficient either, in comparison, tree-based genomes are less vulnerable in this aspect. The loop in Fig. 4.21.2, for instance, stays intact although it is now one useless instruction richer. If we take precautions in LGP in order to mitigate these problems, linear Genetic Programming becomes more competitive to Standard Genetic Programming also in terms of robustness of the recombination operations.

One rather complicated approach to do so are to create intelligent mutation and crossover operators which preserve the control flow of the program when inserting or deleting instructions. Such operations analyze the program structure and automatically correct jump targets in these cases, for instance.

4.6.1 Homologous Crossover

According to Banzhaf et al. [846], natural crossover is very restricted and usually exchanges only genes that express the same functionality and are located at the same position on the chromosomes.

Definition 4.3 (Homology). In genetics, *homology*³⁵ of protein-coding DNA sequences means that they code for the same protein which may indicate common function. Homologous chromosomes³⁶ are chromosomes in a biological cell that pair during meiosis or non-identical chromosomes which code for the same functional feature by containing similar genes in different allelic states.

In other words, homologous genetic material is very similar and in nature, only such material is exchanged in sexual reproduction. In linear Genetic Programming, however, two individuals often differ in their structure and in the function of genes at the same loci. Francone et al. [937] introduce their *sticky crossover* operator which resembles homology by allowing the exchange of instructions between two parents only if they reside at the same loci, too. It first chooses a sequence of code in the first parent and then swaps it with the sequence at exactly the same position in the second parent.

4.6.2 Page-based LGP

A similar approach is the Page-based linear Genetic Programming of Heywood and Zincir-Heywood [1031], where programs are described as sequences of pages, each including the same number of instructions. Here, crossover exchanges only a single page between the parents and, as a result, becomes less destructive.

4.7 Graph-based Approaches

In this section, we will discuss some Genetic Programming approaches that are based on graphs rather than on trees or linear sequences of instructions.

4.7.1 Parallel Algorithm Discovery and Orchestration

Parallel Algorithm Discovery and Orchestration (PADO) is a Genetic Programming method introduced by Teller and Veloso [894] in 1995. In their CEC paper [895] and their 1996 book chapter [896], they describe the graph-structure of their approach as sketched in Figure 4.22. A PADO program is a directed graph of up to n nodes, where each node may have as many as n outgoing arcs which define the possible control flows. A node consists of two parts: *action* and a *branching decision*. The programs used indexed memory and an implicitly accessed stack. The actions pop their inputs from the stack and place their outputs onto it. After a node's action has been executed, the branching decision function is used to determine over which of the outgoing arcs the control will be transferred. It can access the stack, the memory, and the action type of the previous node in order to make that decision.

A program in the PADO-syntax has a start node which will initially receive the control token and an end node which terminates the program after its attached action has been performed. Furthermore, the actions may call functions from a library and automatically defined functions (ADFs). These ADFs basically have same structure as the main program and can also invoke themselves recursively.

³⁵ [http://en.wikipedia.org/wiki/Homology_\(biology\)](http://en.wikipedia.org/wiki/Homology_(biology)) [accessed 2008-06-17]

³⁶ http://en.wikipedia.org/wiki/Homologous_chromosome [accessed 2008-06-17]




Fig. 4.22: The general structure of a Parallel Algorithm Discovery and Orchestration program.

As actions, PADO provides algebraic primitives like $+$, $-$, $*$, $/$, **NOT**, **MAX**, and **MIN**; the memory instructions **READ** and **WRITE**; branching primitives like **IF-THEN-ELSE** and **PIFTE** (alternative with a randomized condition); as well as constants and some domain-specific instructions. Furthermore, actions may invoke more complex library functions or ADFs. An action takes its arguments from the stack and also pushes its results back onto it. The action 6, for instance, pushes 6 on the stack whereas the **WRITE** action pops two values, v_1 and v_2 , from it and pushes the value of the memory cell indexed by v_2 before storing v_1 at this location.

In PADO, so-called SMART operators are used for mutation and recombination which co-evolve with the main population as described in [1032, 896].

4.7.2 Parallel Distributed Genetic Programming

Parallel Distributed Genetic Programming (PDGP) is a method for growing programs in the form of graphs that has been developed by Poli [1033; 1034; 1035; 1036] in the mid 1990s. In PDGP, a graph is represented as a fixed-size, n -dimensional grid. The nodes of the grid are labeled with operations, functions, or references to variables. Except for the latter case, they are connected to their inputs with directed links. Both, the labels as well as the connections in the grid are subject to evolution.




Fig. 4.23.1: tree structure




Fig. 4.23.2: graph structure




Fig. 4.23.3: PDGP structure

Fig. 4.23: The term $\max \{x * y, x * y + 3\}$

In order to illustrate this structure, we use the term $\max\{x * y, x * y + 3\}$ as example. We already have elaborately discussed how we can express mathematical terms as trees. Fig. 4.23.1 illustrates such a function tree. Using a directed graph, as outlined in Fig. 4.23.2, we can retrieve a more compact representation of the same term by reusing the expression $x * y$. Evolving such graphs is the goal of PDGP. Therefore, a grid structure needs to be defined first. In Fig. 4.23.3, we have settled for a two dimensional 4×3 grid. Additionally, we add a row at the top containing one cell for each output of the program. We can easily fill the graph from Fig. 4.23.2 into this grid. This leaves some nodes unoccupied. If we assume that Fig. 4.23.3 represents a solution grown by this Genetic Programming approach, these nodes would be labeled with some unused expressions and would somehow be connected without any influence on the result of the program. Such an arbitrary configuration of *inactive* nodes (or introns and links is sketched in light gray in Fig. 4.23.3. The nodes which have influence on the result of the program, i. e., those which are connected to an output node directly or indirectly, are called *active* nodes.

We may impose restrictions on the connectivity of PDGP graphs. For instance, we can define that each row must only be connected to its predecessor in order to build layered feed-forward networks. We can transform any given parallel distributed program (i. e., any given acyclic graph) into such a layered network if we additionally provide the identity function so pass-through nodes can evolve as shown in Fig. 4.23.3. Furthermore, we could also attach weights to the links between the nodes and make them also subject to evolution. This way, we can also grow artificial neural networks [1037]. However, we can as well do without any form of restrictions for the connectivity and may allow backward connections in the programs, depending on the application.

An interesting part of PDGP is how the programs are executed. Principally, it allows for a great proportion of parallelism. Coming back to the example outlined Fig. 4.23.3, the values of the leaf nodes could be computed in parallel, as well those of the pass-through and the addition node.

Genetic Operations

For this new program representation, novel genetic operations are needed.

Creation

Similar to the *grow* and *full* methods for creating trees in Standard Genetic Programming introduced in Section 4.3.1 on page 132, it is possible to obtain balanced or unbalanced graphs/trees in PDGP, depending whether we allow variables and constants to occur anywhere in the program or only at a given, predetermined depth.

Crossover

SAAN Crossover The basic recombination operation in PDGP is *Sub-graph Active-Active Node (SAAN) crossover*. The idea of SAAN crossover is that active sub-graphs represent functional units which should be combined in different ways in order to explore new, useful constellations. It proceeds as follows:

1. Select a random active node in each parent, the crossover points.
2. Extract the sub-graph that contains all the (active) nodes that influence the result of the node marking the crossover point in the first parent.
3. Insert this sub-graph at the crossover point in the second parent. If its x-coordinate is incompatible and some nodes of the sub-graph would be outside the grid, wrap it so that these nodes are placed on the other side of the offspring.

Of course, we have to ensure that the depths of the crossover points are compatible and no nodes of the sub-graph would “hang” below the grid in the offspring. This can be achieved by first selecting the crossover point in the first parent and then choosing a compatible crossover point in the second parent.

SSAAN Crossover The *Sub-Sub-Graph Active-Active Node (SSAAN) Crossover* method works essentially the same way, with one exception: it disregards crossover point depth compatibility. It may now happen that we want to insert a sub-graph into an offspring at a point where it does not fit because it is too long. Here we make use of the simple fact that the lowest row in a PDGP graph always is filled with variables and constants only – functions cannot occur there because otherwise, no arguments could be connected to them. Hence, we can cut the overhanging nodes of the sub-graph and connect the now unsatisfied arguments at second-to-last row with the nodes in the last row of the second parent. Of course, we have to pay special attention where to cut the sub-graph: terminal nodes that would be copied to the last row of the offspring can remain in it, functions cannot.

SSIAN Sub-Sub-Graph Inactive-Active Node (SSIAN) Crossover works exactly like SSAAN crossover except that the crossover point in the first parent is chosen amongst both, active and inactive nodes.

Mutation

We can extend the mutation operation from Standard Genetic Programming easily to PDGP by creating new, random graphs and insert them at random points into the offspring. In the context of PDGP, this is called *global mutation* and can be achieved by creating a completely new graph and performing crossover with an existing one.

Furthermore, *link mutation* is introduced as an operation that performs simple local changes on the connection topology of the graphs.

ADLs

Similar to Standard Genetic Programming, we can also introduce automatically defined functions³⁷ in PDGP by extending the function set with a new symbol which then executes an (also evolved) subprogram when being evaluated. *Automatically Defined Links*, ADLs, work similarly, except that a link is annotated with the subprogram-invoking symbol [1038, 1039].

4.7.3 Genetic Network Programming

Genetic Network Programming (GNP) is a Genetic Programming technique introduced by Katagiri et al. [1040] at the 2001 CEC conference in Seoul [1040, 1041, 1042, 1043]. In GNP, programs are represented as directed graphs called networks which consist of three types of nodes: the start node, judgment nodes and processing nodes. A processing nodes executes an action from a predefined set of actions P and can have exactly one outgoing connection to a successor node. Judgment nodes may have multiple outgoing connections and have one expression from the set of possible judgment decisions J attached to them with which they make this decision. As in the example illustrated in Figure 4.24, each node in the network is represented by two genes, a node gene and a connection gene. The node gene consists of two values, the node type (which is 0 for processing nodes and 1 for judgment nodes) and the function index. For processing nodes, the function index can take on the values from 0 to $|P| - 1$ and for judgment nodes, it is in $0..|J| - 1$. These values identify the action or decision function to be executed whenever the node receives the control token. In the connection gene, the indices of the other nodes the node is connected to are stored. For processing nodes, this list has exactly one entry, for judgment nodes there always are at least two outgoing connections (in Figure 4.24, there are exactly two). Notice that programs can be interpreted in this representation directly without needing an explicit genotype-phenotype mapping.

Crossover is performed by randomly exchanging notes (and their attached connections) between the parent networks and mutation randomly changes the connections. Murata and

³⁷ More information on ADFs can be found in Section 4.3.9 on page 137.




Fig. 4.24: An example program in Genetic Network Programming syntax.

Nakamura [884] extended their approach in order to evolve programs for multi-agent systems where the behavior of agents depends on the group they are assigned groups to. In this Automatically Defined Groups (ADG) model, an individual is defined as a set of GNP programs [1044, 1045]. Genetic Network Programming has also been combined with reinforcement learning by Mabu et al. [1046].

4.7.4 Cartesian Genetic Programming

Cartesian Genetic Programming (CGP) was developed by Miller and Thomson [1047] in order to achieve a higher degree of effectiveness in learning Boolean functions [1048, 1049]. In his 1999 paper, Miller [1048] explains the idea of Cartesian Genetic Programming with the example of a program with $o = 2$ outputs that computes both, the difference and the sum, of the volumes of two boxes $V_1 = X_1X_2X_3$ and $V_2 = Y_1Y_2Y_3$. As illustrated in Figure 4.25, the $i = 6$ input variables $X_1 \dots X_3$ and $Y_1 \dots Y_3$, placed to the left, are numbered from 0 to 5. As function set, we use $\{+ = 0, - = 1, * = 2, / = 3, \vee = 4, \wedge = 5, \oplus = 6, \neg = 7\}$. Like in PDGP, we define a grid of cells before the evolution begins. In our example, this grid is $n = 3$ cells wide and $m = 2$ cells deep. Each of the cells can accommodate an arbitrary function and has a fixed number of inputs and outputs (in the example $i' = 2$ and $o' = 1$, respectively). The outputs of the cells, similarly to the inputs of the program, are numbered in ascending order beginning with i . The output of the cell in the top-left has number 6, the one of the cell below number 7, and so on. This numeration is annotated in gray in Figure 4.25.

Which functions the cells should carry out and how their inputs and outputs are connected will be decided by the optimization algorithm. Therefore, we could use, for instance, a genetic algorithm with or without crossover or a hill climbing approach. The genotypes of Cartesian Genetic Programming are fixed-length integer strings. They consist of $n * m$ genes, each encoding the configuration of one cell. Such a gene starts with i' numbers identifying the incoming data and one number (underlined in Figure 4.25) denoting the function it will carry out. Another gene at the end of the genotype identifies which of the available data are “wired” to the outputs of the program.

By using a fixed-length genotype, the maximum number of expressions in a Cartesian program is also predefined. It may, however, be shorter, since not all internal cells are




Fig. 4.25: An example for the GPM in Cartesian Genetic Programming.

necessarily connected with the output-producing cells. Furthermore, not all functions need to incorporate all i' inputs into their results. $\neg$, which is also part of the example function set, for instance, uses only the first of its $i' = 2$ input arguments and ignores the second one.

Levels-back, a parameter of CGP, is the number of columns to the left of a given cell whose outputs may be used as inputs of this cell. If levels-back is one, the cell with the output 8 in the example could only use 6 or 7 as inputs. A levels-back value of 2 allows it to be connected with 0-5. Of course, the reproduction operations have to respect the levels-back value set.

CGP labeled itself a Genetic Programming technique from the beginning. However, most of the work contributed about it did not consider a recombination operation. Hence, one could regard it also as an evolutionary programming³⁸ method. Lately, researchers also begin to focus on efficient crossover techniques for CGP [1050].

Neutrality in CGP

Cartesian Genetic Programming explicitly utilizes different forms of neutrality³⁹ in order to foster the evolutionary progress. Normally, neutrality can have positive as well as negative effects on the evolvability of a system. Yu and Miller [179; 1051] outline different forms of neutrality in Cartesian Genetic Programming which also apply to other forms of GP or GAs:

- Inactive genes define cells that are not connected to the outputs in any way and hence cannot influence the output of the program. Mutating these genes therefore has no effect on the fitness and represents *explicit neutrality*.
- Active genes have direct influence on the results of the program. Neutral mutations here are such modifications that have no influence on the fitness. This *implicit neutrality* is the results of functional redundancy or introns.

Their experiments indicate that neutrality can increase the chance of success of Genetic Programming for needle-in-a-haystack fitness landscapes and in digital circuit evolution [178].

³⁸ See Chapter 6 on page 195 for more details on evolutionary programming.

³⁹ See Section 1.4.4 on page 43 and Section 1.4.4 on page 46 for more information on neutrality and redundancy.

Embedded Cartesian Genetic Programming

In 2005, Walker and Miller [1052] published their work on Embedded Cartesian Genetic Programming (ECGP), a new type of CGP with a module acquisition [1053] method in form of automatic module creation [1052, 1054, 1055]. Therefore, three new operations are introduced:

- **Compress** randomly selects two points in the genotype and creates a new module containing all the nodes between these points. The module then replaces these nodes with a cell that invokes it. The compress operator has the effect of shortening the genotype of the parent and of making the nodes in the module immune against the standard mutation operation but does not affect its fitness. Modules are more or less treated like functions so cell to which a module number has been assigned now uses that module as “cell function”.
- **Expand** randomly selects a module and replaces it with the nodes inside. Only the cell which initially replaced the module cells due to the Compress operation can be expanded in order to avoid bloat.
- The new operator **Module Mutation** changes modules by adding or removing inputs and outputs and may also carry out the traditional one-point mutation on the cells of the module.

General Information

Areas Of Application

Some example areas of application of Cartesian Genetic Programming are:

Application	References
Electrical Engineering and Circuit Design	[1056, 178, 1047, 1057, 1058, 1054]
Symbolic Regression and Function Synthesis	[1048, 1049, 179, 1049]
Robotics	[1059, 1060]
Prime Number Prediction	[1061]

Online Resources

Some general, online available resources on Cartesian Genetic Programming are:

http://www.cartesiangp.co.uk/ [accessed 2007-11-02]
Last update: up-to-date
Description: The homepage of Cartesian Genetic Programming
http://www.emoware.org/evolutionary_art.asp [accessed 2007-11-02]
Last update: 2006
Description: A website with art pieces evolved using Cartesian Genetic Programming.

4.8 Epistasis in Genetic Programming

4.8.1 Introduction

In the previous sections, we have discussed many different Genetic Programming approaches like Standard Genetic Programming and the Grammar-guided Genetic Programming family.

We also have elaborated on linear Genetic Programming techniques that encode an algorithm as a stream of instructions, very much like real programs are represented in the memory of a computer.

Whenever we use such methods to evolve real algorithms, we encounter the problem of epistasis which has been introduced earlier in Section 1.4.5 on page 47. In an algorithm, each instruction depends on the instructions that have been executed before. The result of one instruction will influence the behavior of those executed afterwards. Besides this general dependency, we can observe a set of other epistatic effects which we will introduce in this section.

Problems of String-to-Tree GPMs

Many multiple Grammar-guided Genetic Programming methods like Grammatical Evolution, Christiansen Grammar Evolution, and Gads employ a genotype-phenotype mapping between an (integer) string genome and trees that represent sentences in a given grammar. According to Ryan et al. [978], the idea of mapping string genotypes can very well be compared to one of the natural antetypes of artificial embryogeny⁴⁰: the translation of the DNA into proteins. This process depends very much on the proteins already produced and now present around the cellular facilities. If a certain piece of DNA has created a certain protein X and is transcribed again, it may result into a molecule of protein type Y because of the presence of X .

Although this is a nice analogy, it also bears an important weakness. Genomes which exhibit such effects usually suffer from weak causality⁴¹ [810] and only the lord knows why the DNA stuff works at all. In Grammatical Evolution for example, a change in any gene in a genotype g will also change the meaning of all subsequent alleles. This means that mutation and crossover will probably have very destructive impacts on the individuals [260]. Additionally, even the smallest change in the genotype can modify the whole structure and functionality of the phenotype. A valid solution can become infeasible after a single reproduction operation.

Figure 4.26 outlines how the change of a single bit a genotype (in hexadecimal notation) may lead to a drastic modification in the tree structure when a string-to-tree mapping is applied. The resulting phenotype in the example has more or less nothing in common with its parent except maybe the type of its root node.




Fig. 4.26: Epistasis in a Grammatical Evolution-like approach.

⁴⁰ Find out more about artificial embryogeny in Section 3.6 on page 122.

⁴¹ The principles of causality and locality are discussed in Section 1.4.2 on page 40.

The lack of causality is rooted in a strong epistasis in the string-tree-GPM approaches: many loci of the genotypes have mutual influence. The efficiency of the reproduction operations of the mentioned approaches will probably decrease with a growing set of non-terminal symbols and corresponding productions.

One way to decrease these effects is (implicitly) inherent in the Gene Expression Programming approach (see Section 4.4.3 on page 142). An a-priori division of genotypes into different sections, each independently encoding different parts of the phenotype, can reduce the harmful influences of the reproduction operations.

Positional Epistasis

In order to clarify the role of positional epistasis in the context of Genetic Programming, we begin with some basic assumptions. Let us consider a program P as a form of function $P : I \mapsto O$ that connects the possible inputs I of a system to its possible outputs O . Two programs P_1 and P_2 can be considered as equivalent if $P_1(i) = P_2(i) \forall i \in I$.⁴²

For the sake of simplicity, we further define a program as a sequence of n statements $P = (s_1, s_2, \dots, s_n)$. For these n statements, there are $n!$ possible permutations. We argue that the fraction $\theta(P) = \frac{v}{n!}$ of permutations v that leads to programs equivalent to P is a measure of robustness for a given phenotypic representation in Genetic Programming. More precisely, a low value of θ indicates a high degree of epistasis, which means that the loci (the positions) of many different genes in a genome have influence on their functionality [1062]. This reduces, for example, the efficiency of reproduction operations like recombination, since they often change the number and order of instructions in a program. A general rule in evolutionary algorithms is thus to reduce epistasis [264] (see rule 9 in Section 1.5.2 on page 58).

Many of the phenotypic and most genotypic representations in Genetic Programming mentioned so far seem to be rather fragile in terms of insertion and crossover points. One of the causes is that their genomes have high positional epistasis (low θ -measures), as sketched in Figure 4.27.

The points discussed in this section do by no means indicate that the involved Genetic Programming approaches are infeasible or deliver only inferior results. Most of them have provided human-competitive solutions or performed even better. We point out just some classes of problems that, if successfully solved, could even increase their utility.

4.8.2 Algorithmic Chemistry

Algorithmic Chemistries were first discussed by Fontana [1063] on basis of the λ -calculus. The Algorithmic Chemistry approach by Lasarczyk and Banzhaf [1064; 1065], [1066] represents one possible method to circumvent the positional epistasis discussed Section 4.8.1. To put it bluntly, the Algorithmic Chemistries as developed by Lasarczyk and Banzhaf are similar to linear Genetic Programming⁴³ where the execution order of the single instructions is defined by some random distribution instead of being fixed as in normal programs.

This issue can probably be described best using a simple example. Therefore, let us define the set of basic constructs which will make up the programs first. In both, [1066] and [1065], a simple assembler-like language is used where each instruction has three parameters, two input and one output register. Registers are the basic data stores, the variables, of this language and instructions whose results depend only on one input simply ignore their second parameter. In [1066], Banzhaf and Lasarczyk use a language comprising the eight instructions $+$, $-$, $/$, $*$, $\wedge$, **and**, **or**, and **not**. Furthermore, they provided eleven read-only registers with evolved constants and thirty read-write registers, the so-called *connection registers*.

⁴² In order to cover stateful programs, the input set may also comprise sequences of input data.

⁴³ The linear Genetic Programming approach is outlined in Section 4.6 on page 161.




Fig. 4.27.1: In Standard Genetic Programming and symbolic regression




Fig. 4.27.2: In Standard Genetic Programming




Fig. 4.27.3: In linear Genetic Programming




Fig. 4.27.4: With Genotype-Phenotype Mapping, like in Grammatical Evolution-like approaches

Fig. 4.27: Positional epistasis in Genetic Programming.

In Fig. 4.28.1 we have illustrated the sequential structure of a normal program which might have evolved in a linear Genetic Programming experiment. Whenever it is executed, be it for determining its fitness or later, as part of an application, the instructions are processed one after another step by step. This execution scheme is common to all our off-the-shelf PCs where the CPU uses an internal register (the so-called instruction pointer) which points to the instruction to be executed next and which is incremented in this process.

Programs in Algorithmic Chemistry can evolve essentially in the same way as programs in linear Genetic Programming do. As genotypes, similar sequences of instructions can be used. This similarity, however, stops at the phenotypic level⁴⁴. Here, the programs are simple multisets which do not define any order on their elements (the instructions), as sketched in Fig. 4.28.2. When such a program is executed a *random sequencer* draws one element from this set in each time step and executes it.⁴⁵

This approach clearly leads to a θ -value of zero, since all positional dependencies amongst the instructions have vanished. As trade-off, however, there are a number of interesting side

⁴⁴ Although we distinguish between genotype and phenotype, no genotype-phenotype mapping is needed since the randomized execution can perfectly be performed on an array of instructions.

⁴⁵ Banzhaf and Lasarczyk [1066] use the same approach for introducing a new recombination operator which creates an offspring by drawing instructions randomly from both parents.




Fig. 4.28.1: Linear Genetic Programming phenotype and execution.

Fig. 4.28.2: Algorithmic Chemistry phenotype and execution.

Fig. 4.28: The difference between linear Genetic Programming and Algorithmic Chemistries.

effects. Since programs no longer are sequences, there is, for example, no “last instruction” anymore. Thus, the Algorithmic Chemistry programs also have no clear “end”. Therefore, the randomized execution step is performed for a fixed number of steps, five times the number of instructions in [1064]. As pictured in Fig. 4.28.2, a certain instruction may occur multiple times in a program, which increases its probability of being picked for execution. The biggest drawback of this approach is that the programs are no longer deterministic and their behavior and results may vary between two consecutive executions. Therefore, multiple independent runs should always be performed and the median or mean return value of them should be considered as the *true* result. Stripping the instructions of their order also will make it harder for higher-level constructs like alternatives or loops to evolve, let alone modules or functions. On the positive side, it also creates a large potential for parallelization and distribution which could be beneficial especially in multi-processor systems.

4.8.3 Rule-based Genetic Programming

Besides the specialized Algorithmic Chemistry approach of Lasarczyk and Banzhaf [1064; 1065], there exists one very general class of evolutionary algorithms that elegantly circumvents positional epistasis⁴⁶: the Learning Classifier Systems (LCS) family [1067, 1068] which you can find discussed in Chapter 7 on page 197. Let us focus on the Pittsburgh approach associated with Spears and De Jong [1069], where a genetic algorithm evolves a population of rule sets [1070, 1069]. Each individual in this population consists of multiple classifiers (the rules) that transform input signals into output signals. The evaluation order of the rules in such a classifier system C plays absolutely no role except for rules concerning the same output bits, i. e., $\theta(C) \approx 1$.

The basic idea behind the Rule-based Genetic Programming approach is to use this knowledge to create a new program representation that retains high θ -values (and thus, also a higher degree of neutrality) in order to become more robust in terms of reproduction operations [1071]. This will probably lead to a more durable evolution with a higher probability

⁴⁶ You can find positional epistasis discussed in Section 4.8.1 on page 171

of finding good solutions. On the other hand, it also extends LCS by introducing some of the concepts from Genetic Programming like mathematical operations.

This approach to Genetic Programming is illustrated by the example shown in Figure 4.29. Like in Pitt-style Learning Classifier Systems, the programs consist of arbitrary many rules. A rule evaluates the values of some symbols in its condition part (left of $\Rightarrow$) and, in its action part, assigns a new value to one symbol or performs any other procedure specified in its setup. In its structure, the RBGP language is very similar to Dijkstra's Guarded Command Language⁴⁷ (GCL) [1072].⁴⁸




Fig. 4.29: genotype-phenotype mapping in Rule-based Genetic Programming.

Genotype and Phenotype

Before the evolution in Rule-based Genetic Programming begins, the number of symbol and their properties must be specified as well as the possible actions. Each symbol identifies an integer variable, which is either read-only or read-write. Generally, we also define some constants like 0 and 1 and an input symbol **start** which will only be 1 during the first execution of the program and then becomes 0 (but may be written to by the program). Additionally, a program can be provided with some general-purpose variables (a and b in the example). Further symbols with special meanings can be introduced. For evolving distributed algorithms, for instance, an input symbol **receive** where incoming messages

⁴⁷ http://en.wikipedia.org/wiki/Guarded_Command_Language [accessed 2008-07-24]

⁴⁸ I have to thank David R. White for this information – when devising Rule-based Genetic Programming, I didn't even know that the Guarded Command Language existed.

will occur and a variable `send` from which outgoing messages can be transmitted could be added. An action set containing mathematical operations like addition, subtraction, value assignment, and some sort of logical negation ($1-x$) is sufficient in most cases but may arbitrarily be extended. Alternatively to the `send` and `receive` symbol, actions could be defined with the same semantics.

From these specifications, the system can determine how many bits are needed to encode a single rule. The genotypes are bit strings with a length which is a multiple of this number. With this simple genotype, any possible nesting depth of condition statements and all possible logical operations can be encoded. Furthermore, it is even possible to construct a tree-like program structure from the rules, since each of them corresponds to a single `if` statement in a normal programming language.

There are similarities between our RBGP and some special types of LCS, like the abstracted LCS by Browne and Ioannides [1073] and the S-expression-based LCS by Lanzi and Perrucci [1074]. The two most fundamental differences lie in the semantics of both, the rules and the approach: In RBGP, a rule may directly manipulate symbols and invoke external procedures with (at most) two in/out-arguments. This includes mathematical operations like multiplication and division which do not exist a priori in LCS. They would have to evolve on basis of binary operations, which is, although possible, very unlikely. Furthermore, the individuals in RBGP are not classifiers but programs. Classifiers are intended to be executed once for a given situation, judge it, and decide upon an optimal output. A program, on the other hand, runs independently and asynchronously performs a computation and interacts with its environment. Also, the syntax of RBGP is very extensible because the nature of the symbols and actions is not bound to specific data types but can easily be adapted to floating point computation.

Program Execution and Dimensions of Independence

The simplest method to execute a rule-based program would be to loop over it in a cycle. Although this approach is sufficient for simulation purposes, it would result in a large waste of CPU power on a real system. This consumption of computational power can be reduced very much if the conditional parts of the rules are only evaluated when one of the symbols that they access changes.

Positional Independence

Such changes can either be caused by data incoming from the outside, like messages that are received by a distributed algorithm (and are stored in the `receive` symbol in our example) or by the actions invoked by the program itself. In RBGP, actions do not directly modify the values of the symbols but rather write their results to a temporary storage. After all actions have been processed, these values are written back to the real memory, as sketched in Figure 4.30. Therefore, the symbols in the condition part and in the computation parts of the actions are annotated with the index t and those in the assignment part of the actions are marked with $t + 1$.

This approach allows for a greater amount of disarray in the rules, since the only possible positional dependencies left are those of rules that write to the same variables. All other rules can be freely permuted without any influence on the behavior of the program. Hence, the positional epistasis in RBGP is very low.

Cardinality Independence

By excluding any “learning” features like the Bucket Brigade Algorithm⁴⁹ from the evolution, we additionally gain some form of insensitivity in terms of rule cardinality. It is irrelevant whether a rule occurs once, twice, or even more often in a program. If triggered,

⁴⁹ The Bucket Brigade Algorithm is discussed in Section 7.3.7 on page 204.




Fig. 4.30: A program computing the faculty of a natural number in Java and RBGP syntax.

all occurrences of the rule use the same input data and thus will write the same values to the (temporary variable representing) the target symbol. Assuming that an additional objective function which puts pressure into the direction of smaller programs is always imposed, superfluous rules will be wiped out during the course of the evolution anyway.

Neutrality

The existence of neutral reproduction operations can have a positive as well as negative influence on the evolutionary progress (see Section 1.4.4 on page 43). The positional and cardinality independence are a clear example of phenotypic redundancy and neutrality in RBGP. They allow a useful rule to be replicated arbitrary often in the same program without decreasing its functional fitness. This is likely to happen during crossover. By doing so, the conditional parts of the rule will (obviously) be copied too. Subsequent mutation operations may now modify the action part of the rule and lead to improved behavior.

All these new degrees of freedom are achieved without most of the drawbacks that are inherent in Algorithmic Chemistries (which have been discussed in Section 4.8.2). The program flow is fully deterministic and so are its results. Like in Algorithmic Chemistries, it is harder to determine the number of steps needed for program execution, although we can easily detect the termination of local algorithms as the point where an execution does not lead to any change in the symbols.

Complex Statements

From the previous descriptions, it would seem that rule-based programs are strictly sequential, without branching or loops. This is generally not the case.

Complex Conditions

Assume that we have five variables `a...e` and want to express something like

```

1  if( (a<b) ∧ (c>d) ∧ (a<d) ) {
2    a += c;
3    c--; }

```

Listing 4.8: A complex conditional statement.

we can do this in RBGP with 4 rules:

```

1  true  ∧  true  ⇒ et+1 = 0
2  (at < bt) ∧ (ct > dt) ⇒ et+1 = 1
3  (at < dt) ∧ (et = 1) ⇒ at+1 = at + ct
4  (at < dt) ∧ (et = 1) ⇒ ct+1 = ct - 1

```

Listing 4.9: The RBGP version of Listing 4.8.

Although this does not look very elegant, it fulfills the task by storing the first part of the condition as logical value in the variable **e**. **e** will normally be 0 because of line 1 and is only set to 1 by rule 2. Since both of them write to temporary variables, the then-part of the condition (in lines 3 and 4 will be reached in the next round (if **a**<**d** holds too). Notice that the only positional dependency in Listing 4.9 are that rule 2 must always occur after 1 – all rule permutations that obey this statement are equivalent and so is Listing 4.10:

```

1  (at < dt) ∧ (et = 1) ⇒ at+1 = at + ct
2  true  ∧  true  ⇒ et+1 = 0
3  (at < bt) ∧ (ct > dt) ⇒ et+1 = 1
4  (at < dt) ∧ (et = 1) ⇒ ct+1 = ct - 1
5  true  ∧  true  ⇒ et+1 = 0
6  (at < dt) ∧ (et = 1) ⇒ at+1 = at + ct
7  (at < bt) ∧ (ct > dt) ⇒ et+1 = 1
8  (at < dt) ∧ (et = 1) ⇒ ct+1 = ct - 1

```

Listing 4.10: An equivalent alternative version of Listing 4.9.

Loops

Loops in RBGP can be created in the very same fashion.

```

1  b = 1;
2  for(a=c; a>0; a--) {
3    b *= a;
4  }

```

Listing 4.11: A loop in Rule-based Genetic Programming.

The loop defined in Listing 4.11 can be expressed in RBGP as outlined in Listing 4.12, where we use the **start**-symbol (line 1 and 2) to initialize **a** and **b**. As its name suggests, the **start** is only 1 at the very beginning of the program's execution and 0 afterwards (unless modified by an action).

```

1  (startt > 0) ∨ false ⇒ at+1 = ct
2  true  ∧ (startt > 0) ⇒ bt+1 = 1
3  (at > 0) ∧ true  ⇒ at+1 = at - 1
4  false  ∨ (at > 0) ⇒ bt+1 = bt * at

```

Listing 4.12: The RBGP-version of Listing 4.11.

Here no positional or cardinality restrictions occur at all, so Listing 4.13 is equivalent to Listing 4.12.

```

1  false  ∨ (at > 0) ⇒ bt+1 = bt * at
2  true  ∧ (startt > 0) ⇒ bt+1 = 1
3  (startt > 0) ∨ false ⇒ at+1 = ct
4  false  ∨ (at > 0) ⇒ bt+1 = bt * at
5  (at > 0) ∧ true  ⇒ at+1 = at - 1
6  (startt > 0) ∨ false ⇒ at+1 = ct

```

Listing 4.13: An equivalent alternative version of Listing 4.12.

Extended RBGP

We have shown that, although looking rather simple, the primitives of Rule-based Genetic Programming are mighty enough to express many of the constructs known from high-level programming languages. However, in the original RBGP approach, there are some inherent limitations.

The most obvious drawback is the lack of Turing completeness. In order to visualize this problem, imagine the restriction imposed on the Java programming language that only simple types like integer variables and parameters are allowed, but no complex types like arrays. In this case, it would hard to create data structures like lists, since we would have to define a single variable for each and every of their elements. Writing a method for sorting a list of arbitrary length would become impossible. The same goes for Rule-based Genetic Programming as introduced in Section 4.8.3 – the symbols there resemble plain integer variables. In Java, this whole problem is circumvented with arrays, a form of memory which can be accessed indirectly. Thus, adding indirect memory access to the programming language forming the basis of Rule-based Genetic Programming would allow the evolution of list sorting algorithms⁵⁰. We therefore define the notation $[a_t]_t$, which stands for the value of the a_t th symbol (at time step t) in the ordered list of all symbols.

With this extension alone, it becomes possible to use the RBGP language for defining sorting algorithms. Assume that we have the following symbols $(i_0, i_1, \dots, i_{n-1}, l, a, b)$, where the symbols $i_0..i_{n-1}$ constitute the memory which can be used to store the list elements, l is initialized with the length of the list, i.e., the number of the i -elements actually used, and a and b are multi-purpose variables. In the symbol list, i_0 is at position 0, l at position n , a at index $n+1$ and so on. Listing 4.14 performs a form of selection sort⁵¹. Since the modifications performed by RBGP operators are not committed before all rules were applied, no explicit temporary variable is required in the third and fourth rule.

```

1 (startt > 0) ∧ true ⇒ at+1 = 0
2 (startt > 0) ∧ true ⇒ bt+1 = 0
3 (at < lt) ∧ ([at]t < [bt]t) ⇒ [at]t+1 = [bt]t
4 (at < lt) ∧ ([at]t < [bt]t) ⇒ [bt]t+1 = [at]t
5 (bt ≥ at) ∧ (at < lt) ⇒ at+1 = at + 1
6 (bt < at) ∧ true ⇒ bt+1 = bt + 1
7 (bt ≥ at) ∧ (at < lt) ⇒ bt+1 = 0

```

Listing 4.14: A simple selection sort algorithm written in the eRBGP language.

Listing 4.9, one of our previous examples shows another feature of RBGP which might prove troublesome: The condition part of a rule always consists of two single conditions. This is totally unimportant as long as a logical expression to be represented contains only one or two conditions. (If it needs only one, the second condition of the rule may be set to **true** and concatenated with an $\wedge$ -operator.) In Listing 4.9, we try, however, to represent the alternative from Listing 4.8 which consists of three conditions. In order to do so, we needed to introduce the additional symbol **e**. Here we can draw a rough comparison to the human memory⁵² which may be divided into procedural⁵³ (implicit⁵⁴) memory storing, for instance, motor skills and declarative⁵⁵ (explicit⁵⁶) memory holding facts and data. The expressiveness of RBGP's procedural memory is limited, which needs to be mitigated by using more of it and storing additional information in the declarative memory. This is like if we were not able to learn the complete motion of lifting a jar to our lips and instead could

⁵⁰ Well, at least in theory.

⁵¹ http://en.wikipedia.org/wiki/Selection_sort [accessed 2008-05-09]

⁵² <http://en.wikipedia.org/wiki/Memory> [accessed 2008-05-08]

⁵³ http://en.wikipedia.org/wiki/Procedural_memory [accessed 2008-05-08]

⁵⁴ http://en.wikipedia.org/wiki/Implicit_memory [accessed 2008-05-08]

⁵⁵ http://en.wikipedia.org/wiki/Declarative_memory [accessed 2008-05-08]

⁵⁶ http://en.wikipedia.org/wiki/Explicit_memory [accessed 2008-05-08]

only learn how to lift a jar from the table and how to move a (already lifted) jar to our mouth while needing to explicitly remember that both moves belong together. Well, this comparison may be farfetched, but it illustrates that Rule-based Genetic Programming may need to go through a seemingly complex learning process for building a simple algorithm. We therefore extend its expressiveness dropping the constraints on the structure of its rules which increases the number of ways that RBGP can utilize for representing complex rules. This will cost us the ability of using genetic algorithms with fixed-size genes for its evolution and may bring back some positional epistasis which we had previously successfully decreased, though.

```

1 ((at < bt) ∧ ((ct > dt) ∧ (at < dt))) ⇒ at+1 = (at + ct)
2 ((at < bt) ∧ ((ct > dt) ∧ (at < dt))) ⇒ ct+1 = (ct - 1)

```

Listing 4.15: The eRBGP version of Listing 4.8 and Listing 4.9.

```

1 (startt ≠ 0) ⇒ bt+1 = 1
2 (startt ≠ 0) ⇒ ct+1 = at
3 (at > 0) ⇒ at+1 = (at - 1)
4 (at > 0) ⇒ bt+1 = (bt * at)

```

Listing 4.16: The eRBGP version of Listing 4.11 and Listing 4.12.

In Listing 4.15 and Listing 4.16, we repeat the RBGP examples Listing 4.9 and Listing 4.12 – this time in eRBGP syntax. As already mentioned, we cannot use a simple genetic algorithm to evolve these programs, since their structure does not map to a fixed gene size anymore. However, tree-based Standard Genetic Programming as discussed in Section 4.3 can perfectly fulfill this purpose. Listing 4.15, for instance, fits to the tree phenotype depicted in Figure 4.31.




Fig. 4.31: The tree phenotype (and genotype) of Listing 4.15.

With these changes, Extended RBGP becomes much more powerful in comparison with Rule-based Genetic Programming and able to evolve complicated algorithms and data structures. Matter of fact, the proof for Turing completeness of Genetic Programming languages with indexed memory by Teller [1076] could be easily adapted to Extended RBGP.

4.8.4 Soft Assignment

One of these methods is the soft assignment idea (*memory with memory* of McPhee and Poli [1077]). It implicitly targets semantic epistasis by weakening the variable assignment. In

traditional programs, instructions like $x=y$ or `mov x, y` will completely overwrite the value of x with the value of y . McPhee and Poli replace this strict assignment's semantics with

$$x_{t+1} = y_t \equiv x_{t+1} \leftarrow \gamma y_t + (\gamma - 1)x_t \quad (4.3)$$

where x_{t+1} is the value that the variable x will have after and x_t its value before the assignment. y_t is the value of an arbitrary expression which is to be stored in x . The parameter γ is “a constant that indicates the *assignment hardness*” [1077]. For $\gamma = 1$, we have the normal assignment as used in normal programming, for $\gamma = 0$, the values of the variables cannot be changed. McPhee and Poli [1077] report that $\gamma = 0.7$ performed best (better than $\gamma = 1$) when applied to different symbolic regression⁵⁷ problems. It will be interesting to see whether this approach can also be used in general algorithms or is restricted to mathematical problems like function approximation.

4.9 Artificial Life and Artificial Chemistry

It is not hard to imagine what artificial life is. Matter of fact, I assume that everyone of us has already seen numerous visualizations and simulations showing artificial creatures. Even some examples from this book like the Artificial Ant may well be counted to that category.

Definition 4.4 (Artificial Life). Artificial life⁵⁸, also abbreviated with *ALife* or *AL*, is a field of research that studies the general properties of life by synthesizing and analyzing life-like behavior [1078].

Definition 4.5 (Artificial Chemistry). The area of artificial chemistries⁵⁹ subsumes all computational systems which are based on simulations of entities similar to molecules and the reactions amongst them.

According to Dittrich et al. [1079], an artificial chemistry is defined by a triple (S, R, A) , where $S = \{s_1, s_2, s_3, \dots\}$ is the set of possible molecules S , R is the set of reactions that can occur between them, and A is an algorithm defining how the reactions are applied.

Artificial chemistry and artificial life strongly influence each other and often merge into each other. The work of Hutton [1080; 1081; 1082], for example, focuses on generating and evolving self-replicating molecules and cells. There also exists real-world applications of artificial chemistry in many areas of chemistry, computer networking, economics, and sociology. The Algorithmic Chemistries, which we have analyzed in Section 4.8.2 on page 171, are also closely related to artificial chemistries, for instance. In this section, we will discuss some more artificial life and artificial chemistry approaches in the context of Genetic Programming.

4.9.1 Push, PushGP, and Pushpop

In 1996, early research in self-modification or self-evolution of programs has been conducted by Spector and Stoffel [1083] in form of the *ontogenic* extension of their HiGP system [1084]. Basically, they extended the programs evolved with linear Genetic Programming method with the capabilities of shifting and copying segments of their code at runtime.

About half of a decade later, Spector [1085] developed *Push*, a stack-based programming language especially suitable for evolutionary computation [1085, 1086, 933]. Programs in that language can be evolved by adapting existing Standard Genetic Programming systems (as done in *PushGP*) or, more interestingly, by themselves in an autoconstructive manner, which has been realized in the *Pushpop* system. Currently, the Push language is currently available in its third release, *Push3* [1087, 1088].

A Push program is either a single instruction, a literal, or a sequence of zero or more Push programs inside parentheses.

⁵⁷ See Section 22.1 on page 341 for more information on symbolic regression.

⁵⁸ http://en.wikipedia.org/wiki/Artificial_life [accessed 2007-12-13]

⁵⁹ http://en.wikipedia.org/wiki/Artificial_chemistry [accessed 2008-05-01]

```
1 program ::= instruction | literal | ( {program} )
```

An instruction may take zero or more arguments from the stack. If insufficient many arguments are available, it acts as NOOP, i.e., does nothing. The same goes if the arguments are invalid, like when a division by zero would occur.

In Push, there is a stack for each data type, including integers, Boolean values, floats, name literals, and code itself. The instructions are usually named according to the scheme `<type>.<operation>`, like `INTEGER.+`, `BOOLEAN.DUP`, and so on. One simple example for a Push program borrowed from Spector [933], Spector et al. [1087] is

```
1 ( 5 1.23 INTEGER.+ ( 4 ) INTEGER.- 5.67 FLOAT.* )
2 Which will leave the stacks in the following states:
3 FLOAT STACK : (6.9741)
4 CODE STACK : ( ( 5 1.23 INTEGER.+ ( 4 ) INTEGER.- 5.67
5               FLOAT.* ) )
6 INTEGER STACK: (1)
```

Listing 4.17: A first, simple example for a Push program.

Since all operations take their arguments from the corresponding stacks, the initial `INTEGER.+` does nothing because only one integer, 5, is available on the `INTEGER` stack. `INTEGER.-` subtracts the value on the top of `INTEGER` stack (4) from the one beneath it (5) and leaves the result (1) there. On the float stack, the result of the multiplication `FLOAT.*` of 1.23 and 5.67 is left while the whole program itself resides on the `CODE` stack.

Code Manipulation

One of the most interesting features of Push is that we can easily express new forms of control flow or self-modifying code with it. Here, the `CODE` stack and, since Push3, the `EXEC` stack play an important role. Let us take the following example from [1085, 1087]:

```
1 (CODE.QUOTE (2 3 INTEGER.+) CODE.DO)
```

Listing 4.18: An example for the usage of the `CODE` stack.

The instruction `CODE.QUOTE` leads to the next piece of code ((2 3 `INTEGER.+`) in this case) being pushed onto the `CODE` stack. `CODE.DO` then invokes the interpreter on whatever is on the top of this stack. Hence, 2 and 3 will land on the `INTEGER` stack as arguments for the following addition. In other words, Listing 4.18 is just a complicated way to add $2 + 3 = 5$.

```
1 (CODE.QUOTE
2   (CODE.QUOTE (INTEGER.POP 1)
3     CODE.QUOTE (CODE.DUP INTEGER.DUP 1 INTEGER.- CODE.DO INTEGER.*)
4     INTEGER.DUP 2 INTEGER.< CODE.IF)
5 CODE.DO)
```

Listing 4.19: Another example for the usage of the `CODE` stack.

Listing 4.19 outlines a Push program using a similar mechanism to compute the factorial of an input provided on the `INTEGER` stack. It first places the whole program on the `CODE` stack and executes it (with the `CODE.DO` at its end). This in turn leads on the code in lines 2 and 3 being placed on the code stack. The `INTEGER.DUP` instruction now duplicates the top of the `INTEGER` stack. Then, 2 is pushed and the following `INTEGER.<` performs a comparison of the top two elements on the `INTEGER` stack, leaving the result (`true` or `false`) on the `BOOLEAN` stack. The instruction `CODE.IF` executes one of the top two items of the `CODE` stack, depending on the value it finds on the `BOOLEAN` stack and removes all three elements. So in case that the input element was smaller than 2, the top element of the `INTEGER` stack will be removed and 1 will be pushed into its place. Otherwise, the next instruction `CODE.DUP` duplicates the whole program on the `CODE` stack (remember, that everything else has already been removed from the stack when `CODE.IF` was executed). `INTEGER.DUP` copies the

top of the `INTEGER` stack, 1 is stored and then subtracted from this duplicate. The result is then multiplied with the original value, leaving the product on the stack. So, Listing 4.19 realizes a recursive method to compute the factorial of a given number.

Name Binding

As previously mentioned, there is also a `NAME` stack in the Push language. It enables us to bind arbitrary constructs to names, allowing for the creation of named procedures and variables.

```
1 ( DOUBLE CODE.QUOTE ( INTEGER.DUP INTEGER.+ ) CODE.DEFINE )
```

Listing 4.20: An example for the creation of procedures.

In Listing 4.20, we first define the literal `DOUBLE` which will be pushed onto the `NAME` stack. This definition is followed by the instruction `CODE.QUOTE`, which will place code for adding an integer number to itself on the `CODE` stack. This code is then assigned to the name on top of the `NAME` stack (`DOUBLE` in our case) by `CODE.DEFINE`. From there on, `DOUBLE` can be used as a procedure.

The EXEC Stack

Many control flow constructs of Push programs up to version 2 of the language are executed by similar statements in the interpreter. Beginning with Push3, all instructions are pushed onto the new `EXEC` stack prior their invocation. Now, now additional state information or flags are required in the interpreter except from the stacks and name bindings. Furthermore, the `EXEC` stack supports similar manipulation mechanisms like the `CODE` stack.

```
1 ( DOUBLE EXEC.DEFINE ( INTEGER.DUP INTEGER.+ ) )
```

Listing 4.21: An example for the creation of procedures similar to Listing 4.20.

The `EXEC` stack is very similar to the `CODE` stack, except that its elements are pushed in the inverse order. The program in Listing 4.21 is similar to Listing 4.20 [1087].

Autoconstructive Evolution

Push3 programs can be considered as tree structures and hence be evolved using standard Genetic Programming. This approach has been exercised with the *PushGP* system [1085, 1086, 882, 1089, 1090]. However, the programs can also be equipped with the means to create their own offspring. This idea has been realized in a software called *Pushpop* [1085, 1086, 1091]. In Pushpop, whatever is left on top of the `CODE` stack after a programs execution is regarded as its child. Programs may use the above mentioned code manipulation facilities to create their descendants and can also access a variety of additional functions, like

- `CODE.RAND` pushes newly created random code onto the `CODE` stack.
- `NEIGHBOR` takes an integer n and returns the code of the individual in distance n . The population is defined as a linear list where siblings are grouped together.
- `ELDER` performs a tournament between n individuals of the previous generation and returns the winner.
- `OTHER` performs a tournament between n individuals of the current generation, comparing individuals according to their parents fitness, and returns the winner.

After the first individuals able to reproduce have been evolved the system can be used to derive programs solving a given problem. The only external influence on the system is a selection mechanism required to prevent uncontrolled growth of the population by allowing only the children of fit parents to survive.

4.9.2 Fraglets

In his seminal work, Tschudin [1092] introduced Fraglets⁶⁰, a new artificial chemistry suitable for the development and even evolution of network protocols. Fraglets represent an execution model for communication protocols which resembles the chemical reactions in living organisms.

How do Fraglets work?

From the theoretical point of view, the Fraglet approach is an instance of Post's string rewriting systems⁶¹ [1093]. Fraglets are symbolic strings of the form $[s_1 : s_2 : \dots : s_n]$ that represent control information and payload. Each node in the network has a Fraglet store which corresponds to a reaction vessel in chemistry. It is only natural that reaction vessels may contain equal molecules multiple times. The same goes for Fraglet stores which can be implemented as multisets keeping track on the multiplicity of the Fraglets they contain.

Tschudin [1092] defines a simple prefix programming language for Fraglets. It has a fixed instruction set comprising transformation and reaction rules. Transformations like *dup* and *nop* modify a single Fraglet whereas the reactions *match* and *matchP* combine two Fraglets. For the definition of these rules in Table 4.2, we will use the syntax $[s_1 : s_2 : \dots : tail]$ where s_i is a symbol and *tail* is a possibly empty sequence of symbols.⁶²

Examples

After defining the basics of the Fraglet approach, let us now take a look on a few simple examples.

Election

Election in a distributed system means to select one node in the network and to ensure that all nodes receive knowledge of the ID of the selected one. One way to perform such an election is to determine the maximum ID of all nodes, which is what we will do here.

First of all, we define five additional symbols *A*, *B*, *C*, *L*, and *R*. These symbols do not react on their own behalf, i. e., $[A : tail] \longrightarrow [A : tail]$. After a bootstrap reaction, each node in the network will contain a Fraglet of the form $[L : \{id\}]$ containing the identifier *id* of the node that it thinks has won/currently leads in the election. It broadcasts this information to its neighbors, which will receive it in the form of the Fraglet $[R : \{id\}]$. *A*, *B*, and *C* have no further meaning. The election algorithm is constituted by six Fraglets $[node : L]$ which creates the first *L*-type Fraglet at bootstrap, $[matchP : L : fork : L : A]$ and $[matchP : A : broadcast : R]$ which are used to transmit the highest node ID currently known in a *R*-type Fraglet, $[matchP : R : match : L : lt : B : C]$ which initiates a comparison with incoming Fraglets of that type, and $[matchP : B : pop2 : null : L]$ and $[matchP : B : pop2 : L : null]$ evaluating the outcome of this comparison and producing a new *L*-type Fraglet. Figure 4.32 shows the flow of reactions in this system, which will finally lead to all nodes knowing the highest node ID in the network.

⁶⁰ See <http://en.wikipedia.org/wiki/Fraglets> [accessed 2008-05-02] and <http://www.fraglets.net/> [accessed 2008-05-02] for more information.

⁶¹ http://en.wikipedia.org/wiki/Post_canonical_system [accessed 2008-05-02]

⁶² See <http://www.fraglets.net/frag-instrset-20070924.txt> [accessed 2008-05-02] for the full instruction set as of 2007-09-24.

tag	transformation/reaction
basic transformations	
dup	$[dup : t : a : tail] \longrightarrow [t : a : a : tail]$ duplicate a single symbol
exch	$[exch : t : a : b : tail] \longrightarrow [t : b : a : tail]$ swap two tags
fork	$[fork : a : b : tail] \longrightarrow [a : tail], [b : tail]$ copy fraglet and prepend different header symbols
nop	$[nop : tail] \longrightarrow [tail]$ does nothing (except consuming the instruction tag)
null	$[null : tail] \longrightarrow \emptyset$ destroy a fraglet
pop2	$[pop2 : h : t : a : tail] \longrightarrow [h : a], [t : tail]$ pop head element a out of a list $[a : b : tail]$
split	$[split : tail_1 : * : tail_2] \longrightarrow [tail_1], [tail_2]$ break a fraglet into two at the first occurrence of $*$
arithmetic transformations	
sum	$[sum : t : m : n : tail] \longrightarrow [t : m + n : tail]$ a logic operation comparing two numbers
lt	$[lt : yes : no : a : b : tail] \longrightarrow \begin{cases} [yes : \{a\} : \{b\} : tail] & \text{if } a < b \\ [no : \{a\} : \{b\} : tail] & \text{otherwise} \end{cases}$ a logic operation comparing two numbers a and b
communication primitives	
broadcast	$[broadcast : tail] \longrightarrow_n [tail]$ broadcast $tail$ to all nodes n in the network N
send	$[send : dest : tail] \longrightarrow_{dest} [tail]$ send $tail$ to a single node $dest$
node	$_n [node : t : tail] \longrightarrow_n [t : \{id(n)\} : tail]$ obtain the ID $id(n)$ of the current node n
reactions	
match	$[match : a : tail_1], [a : tail_2] \longrightarrow [tail_1 : tail_2]$ two fraglets react, their tails are concatenated
matchP	$[matchP : a : tail_1], [a : tail_2] \longrightarrow [matchP : a : tail_1], [tail_1 : tail_2]$ “catalytic match”, i. e., the matchp rule persists

Table 4.2: Some Fraglet instructions (from [1092] and <http://www.fraglets.net/>).

Quines

Definition 4.6 (Quine). A *quine*⁶³ is a computer program which produces a copy of itself (or its source code) as output.

From Kleene’s second recursion theorem⁶⁴ [1094], it follows that quines can be defined in each Turing complete language. Yamamoto et al. [1095] have introduced quine Fraglets like the one in Figure 4.33 as vehicle for self-replicating and self-modifying programs.

According to Tschudin [1092], Fraglets are predestined for automated protocol synthesis via evolutionary algorithms. In Section 23.1.3 on page 361, we discuss the application of Fraglets for protocol evolution based on his work and in Section 23.1.3 on page 362, the joint work of Yamamoto and Tschudin [707] on online adaptation of Fraglet protocols is outlined.

⁶³ http://en.wikipedia.org/wiki/Quine_%28computing%29 [accessed 2008-05-04]

⁶⁴ http://en.wikipedia.org/wiki/Kleene%27s_recursion_theorem [accessed 2008-05-04]




Fig. 4.32: A Fraglet-based election algorithm




Fig. 4.33: A simple quine Fraglet (borrowed from [1095])

4.10 Evolving Algorithms

With Genetic Programming, we can breed algorithms and programs suitable for a solving given problem classes. In order to guide such an evolutionary process, we have to evaluate the utility of these grown programs. Algorithms are assessed in terms of functional and non-functional requirements. The functional properties comprise all features regarding how good the algorithm solves the specified problem and the non-functional aspects are concerned with, for example, its size and memory consumption. We specify (multiple) objective functions $F = \{f_1, \dots, f_n\}$ to map these attributes to the subsets $Y_1, \dots, Y_n$ of the real numbers $\mathbb{R}$.

4.10.1 Restricting Problems

While some of the non-functional objective values can easily be computed (for example the algorithms size), determining the functional utility of an algorithm requires training. The functional utility of the algorithm specifies how close the results it produces from given inputs come to the output values wanted. This proximity cannot be computed by any arithmetically closed function or algorithm since it is an instance of the Entscheidungsproblem⁶⁵ [1096] as well as of the Halting Problem⁶⁶ [1097].

Entscheidungsproblem

The *Entscheidungsproblem* formulated by David Hilbert [1098, 1099] asks for an algorithm that, if provided with a description of a formal language and statement in that language, can decide whether the statement is **true** or **false**. In the case of Genetic Programming, the formal language is the language in which the algorithms are evolved and the statements are the algorithms itself. Church [1100; 1101] and Turing [1075; 1102] both proved that such an algorithm cannot exist.

Halting Problem

The *Halting Problem* asks for an algorithm that decides if another algorithm will finish at some time or runs forever if provided with a certain, finite input. Again, Turing [1075; 1102] proved that an algorithm solving the Halting Problem cannot exist. His proof is based on a counter-example. If we assume that a correct algorithm `doesHalt` exists (as presumed in Algorithm 4.2) which takes an algorithm as input and determines if the algorithm will halt or not. We could now specify the algorithm `trouble` which in turn uses `doesHalt` to determine if it itself will halt at some point of time. If `doesHalt` returns **true**, it loops forever. Otherwise it halts. `doesHalt` cannot return the correct result for `trouble` and, hence, it cannot be universally applied and thus it is not possible to solve the halting problem algorithmically.

In order to circumvent the Halting Problem, additional measures have to be taken into consideration. In linear Genetic Programming it is for instance easily possible to limit the runtime of programs by defining an upper bound for the number of instructions being executed. This is harder for tree-based representation. We could also prohibit the evolution of infinite loops from the start by limiting the structural elements in the programming language. If there are no loops, there surely cannot be infinite ones either. Teller [1103] suggests further time-limiting approaches but also the use of so-called *anytime algorithms*, i. e., algorithms that store their best-guess of the result in a certain memory cell and update it during their run. They can be stopped at any time, since the result is always there although it may have been refined in the future course of the algorithm.

⁶⁵ <http://en.wikipedia.org/wiki/Entscheidungsproblem> [accessed 2007-07-03]

⁶⁶ http://en.wikipedia.org/wiki/Halting_problem [accessed 2007-07-03]

Algorithm 4.2: *Halting Problem: reductio ad absurdum*

```

1 begin
2   doesHalt(algo) ∈ {true, false}
3   begin
4     | ...
5   end
6   Subalgorithm trouble()
7   begin
8     if doesHalt(trouble) then
9       | while true do
10        | | ...
11    end
12 end

```

4.10.2 Why No Exhaustive Testing?

Since we cannot use some kind *algorithms* in order to determine the evolved programs will terminate, let alone if they will provide correct results, we must use training cases in order to test whether they are suitable for a given problem during.

Software testing is a very important field in software engineering [1104, 1105, 1106, 1107]. The core problem of testing programs for their functionality and performance is the size of the input space. Let us assume that we have a program that takes two integer numbers (32 bit) as input and computes another one as output. For testing this program for all possible inputs, $2^{32} * 2^{32} = 2^{64} = 18\,446\,744\,073\,709\,551\,616$ single tests would be required. Even if each test run only took 1μs, exhaustive testing would take approximately 584 542 years. In most practical applications, the input space is much larger and in general, we can consider it to be countable infinite large according to Definition 26.7 on page 410. Exhaustive testing the evolved algorithms is thus not feasible Teller [1103].

Thus, we can only pick a very small fraction of the possible test scenarios for training and hope that they will provide significant results. The probability that this will happen depends very much on the method with which we select the training cases. Here we can draw another interesting analogy to nature: One can never be sure whether evolved behavioral patterns (like algorithms) are perfect and free from error in all possible situations. Nature here indeed has the same problem as the noble scientists who try to apply Genetic Programming.

What spontaneously comes to my mind here is a scheme in the behavior of monkeys. If a monkey sees or smells something delicious in, for example, a crevice, it sticks its hand in, grabs the item of desire and pulls it out. This simple behavior itself is quite optimal and has served the monkeys well for generations. But with the occurrence of the hairless apes called *homo sapiens*⁶⁷, the situation changed. African hunters still use this behavior against the monkey by creating a situation that was never relevant during its evolutionary *testing* period. They slice a coconut in half and put a hole in one side just big enough for a monkey hand to fit through. Now they place an orange between the two coconut halves, tie them closely together and secure the trap with a rope to a tree. Sooner or later a monkey will smell a free lunch, find the coconut with the hole, stick its hand inside, and grab the fruit. However, with the orange in its fist, it cannot pull the hand out anymore. The hunter can now easily catch the monkey, to whom it never occurs that it could let go the fruit and save its life. So, although evolutionary algorithms like Genetic Programming may provide good solutions for many problems, their results still need to be analyzed and interpreted by somebody at least a bit more cunning than an average monkey⁶⁸.

⁶⁷ http://en.wikipedia.org/wiki/Homo_sapiens [accessed 2008-03-25]

⁶⁸ Use this as inspiration to diss the person sitting next to you.

From the impossibility of deciding the Halting Problem, another important fact follows. The algorithms evolved could, for instance, be compiled to real machine code and executed in order to test them. It is then, however, possible that these processes will never terminate. We therefore either need a measure to restrict the runtime of the executed programs which might be more or less difficult, depending on the operating system we use, or we could use simulations. The latter approach is more suitable. It is very easy to limit the runtime of a program that is executed on a simulator – even to an exact number of CPU ticks, which, in general, would not be possible with real processes. Furthermore, the simulation allows us a much deeper insight into the behavior of the algorithm and we may also emulate resources required like network connections, sensors, or physical phenomena like interference.

4.10.3 All-Or-Nothing?

Even if we can determine more or less exactly whether algorithms work correctly or not, the evolution of them often proves as a special instance of the needle-in-a-haystack problem (see Section 1.4.4 on page 47). From a naïve (and at the same time, mathematically precise) point of view, an algorithm computing the greatest common divisor of two numbers either is correct or wrong. Approaching this problem straightforwardly, we have a Boolean objective function which takes on two values. This issue is called the *all-or-nothing* problem in Genetic Programming. In such a fitness landscape, a few steep spikes representing the correct algorithms are distributed over a large plane of infeasible solution candidates with zero fitness.

The negative influence of all-or-nothing problems have been reported from many areas of Genetic Programming such as the evolution of distributed protocols [1092] (see Section 23.1.3), quantum algorithms [933], expression parsers [1108], and mathematical algorithms (such as the GCD).

In Section 20.3.2, we show how to some means to mitigate this problem for the GCD evolution. However, like those mentioned in some of the previously cited works, such methods are normally application dependent and often cannot be transferred to other problems in a simple manner.

4.10.4 Non-Functional Features of Algorithms

Besides evaluating an algorithm in terms of its functionality, there always exists a set of non-functional features that should be regarded too.

Code Size

In Section 29.1.1 on page 479, we define what algorithms are: compositions of atomic instructions that, if executed, solve some kind of problem or a class of problems. Without specifying any closer what *atomic instructions* are, we can define the following:

Definition 4.7 (Code Size). The code size of an algorithm or program is the number of atomic instructions it is composed of.

The atomic instructions cannot be broken down into smaller pieces. Therefore, the code size is a positive integer number in $\mathbb{N}_0$. Since algorithms are statically finite per definition (see Definition 29.9 on page 482), the code size is always finite.

Code Bloat

Definition 4.8 (Bloat). In Genetic Programming, *bloat* is the uncontrolled growth in size of the individuals during the course of the evolution [1109, 1110, 846, 65, 1111].

The term code bloat is often used in conjunction with code *introns*, which are regions inside programs that do not contribute to the fitness because they can never be reached (see Definition 3.2 on page 118). Limiting the code size and increasing the code efficiency by reducing the number of introns is an important task in Genetic Programming since disproportionate program growth has many bad side effects like:

1. Programs become unnecessarily big – solutions for a problem should always be as small as possible.
2. Mutation and recombination operators always have to select the point in an individual where they will apply their changes. If there are many points in an individual that do not contribute to its functionality, the probability of selection such a point for modification is high. The modified offspring will then have exactly the same functionality as its parents and the genetic operation performed was literally useless.
3. Bloat slows down both, the evaluation [95] and the breeding process of new solution candidates.
4. It leads to increased memory consumption of the Genetic Programming system.

There are many theories about how code bloat emerges [1109], some of them are:

- Unnecessary code hitchhikes with good individuals. Since it is part of a fit solution candidate that creates many offspring, it is likely to be part of many new individuals. According to Tackett [1112], high selection pressure is thus likely to cause code growth. This idea is supported by the research of Langdon and Poli [1111], Smith and Harries [1113], and Gustafson et al. [95].
- As already stated, unnecessary code makes it harder for genetic operations to alter the functionality of an individual. In most cases, genetic operators yield offspring with worse fitness than its parents. If a good solution candidate has good objective values, unnecessary code can be one way of defense against recombination and mutation. If the genetic operators are neutralized, the offspring will have the same good fitness as its parent. This idea has been suggested in many sources, such as [1110, 934, 1114, 1115, 1116, 846, 1117, 1113]. From this point of view, introns are a “bad” form of neutrality⁶⁹. The reduction of the destructive effect of recombination on the fitness may also have positive effects, as pointed out by Nordin et al. [1030; 1118], since it may lead to a more durable evolution.
- Similar to the last theory, the idea of removal bias by Soule and Foster [1119] states that removing code from an individual will preserve the individual’s functionality if the code removed is non-functional. Since the portion of useless code inside a program is finite, there also exists an upper limit of the amount of code that can be removed without altering the functionality of the program. For the size of new sub-trees that could be inserted due to mutation or crossover, on the other hand, no such limit exists. Therefore, programs tend to grow [1119, 1117].
- According to the diffusion theory of Langdon et al. [1117], the number of large programs in the problem space that are functional is higher than the number of small functional programs. Thus, code bloat would be the movement of the population into this direction.
- Another theory states that there exist invalidators that make code unreachable or dysfunctional. In the formula $4+0*(4-x)$ for example, 0 invalidates the whole part $0*(4-x)$. Luke [1109] argues that the chance of the occurrence of inhibitors close to the root of a tree would be higher if the proportion of them remained constant. Hence, their influence on large trees would be bigger than on small trees and with the growth of the programs the amount of unnecessary instructions would increase.
- Luke [1109] furthermore defines a more general theory for the tree growth. As already stated, recombination is likely to destroy the functionality of an individual. On the other hand, the deeper the crossover occurs in the tree, the lesser is its influence because fewer

⁶⁹ You can find the topic of neutrality discussed in Section 1.4.4 on page 43.

instructions are removed. On the other hand, if only a few instructions are replaced from a functional program, they are likely to be exchanged by a larger sub-tree. A new offspring that retains the functionality of its parents therefore tends to be larger.

- Instead of being real solutions, programs that grow uncontrolled also tend to be some sort of decision tables. This phenomenon is called *overfitting* and is discussed in Section 1.4.6 on page 48 and Section 22.1.3 on page 343 in detail. The problem is that overfitted programs tend to have marvelous good fitness for the training cases/sample data, but are normally useless for any other input.
- Gustafson et al. [95] link code growth to loss of diversity.

Some approaches for fighting bloat are discussed in Section 4.10.4.

Runtime and Memory Consumption

Another aspect subject to minimization is generally the runtime of the algorithms grown. The amount of steps needed to solve a given task, i. e., the time complexity, is only loosely related to the code size. Although large programs with many instructions tend to run longer than small programs with few instructions, the existence of loops and recursion invalidates a direct relation.

Like the complexity in time, the complexity in memory space of the evolved solutions often is minimized, too. The number of variables and memory cells needed by program in order to perform its work should be as small as possible. Section 29.1.3 on page 482 provides some additional definitions and discussion about the complexity of algorithms.

Errors

An example for an application where the non-functional errors that can occur should be minimized is symbolic regression. Therefore, the property of *closure* specified in Definition 4.1 on page 148 is usually ensured. Then, the division operator `div` is re-defined in order to prevent division-by-zero errors. Therefore, such a division could either be rendered to a `nop` (i. e., does nothing) or yields 1 or the dividend as result. However, the number of such arithmetical errors could also be counted and made the subject to minimization too.

Transmission Count

If evolving distributed algorithms, the number of messages required to solve a problem should be as low as possible since transmissions are especially costly and time-consuming operations.

Optimizing Non-Functional Aspects

Optimizing the non-functional aspects of the individuals evolved is a topic of scientific interest.

- One of the simplest means of doing so is to define additional objective functions and to perform a multi-objective optimization. Successful and promising experiments by Bleuler et al. [1120], de Jong et al. [100], and Ekárt and Németh [1121] showed that this is a viable countermeasure for code bloat, for instance.
- Another method is limiting the aspect of choice. A very simple measure to limit code bloat, for example, is to prohibit the evolution of trees with a depth surpassing a certain limit [1122].
- Poli [1123] suggests that the fitness of a certain portion of the population with above-average code size should simply be set to the worst possible value. These artificial *fitness holes* will repel the individuals from becoming too large and hence reduce the code bloat.

Evolution Strategy

5.1 Introduction

Evolution Strategies¹ (ES) introduced by Rechenberg [1124; 105; 106] are a heuristic optimization technique based in the ideas of adaptation and evolution, a special form of evolutionary algorithms [1124, 105, 106, 1125, 1126, 440, 1127, 1]. Evolution Strategies have the following features:

- They usually use vectors of real numbers as solution candidates, i. e., $\mathbb{G} = \mathbb{X} = \mathbb{R}^n$. In other words, both the search and the problem space are fixed-length strings of floating point numbers, similar to the real-encoded genetic algorithms mentioned in Section 3.3 on page 117.
- Mutation and selection are the primary operators and recombination is less common.
- Mutation most often changes the elements $\mathbf{x}[i]$ of the solution candidate vector $\mathbf{x}$ to a number drawn from a normal distribution $N(\mathbf{x}[i], \sigma_i^2)$. For reference, you can check Equation 9.1 on page 213 in the text about Random Optimization.
- Then, the values σ_i are governed by self-adaptation [1128, 1129, 608] such as covariance matrix adaptation [1130, 1131, 1132, 1133].
- In all other aspects, they perform exactly like basic evolutionary algorithms as defined in Algorithm 2.1 on page 72.

5.2 General Information

5.2.1 Areas Of Application

Some example areas of application of Evolution Strategy are:

Application	References
Data Mining and Data Analysisanalysis	[464]
Scheduling	[1134]
Chemistry, Chemical Engineering	[1135, 1136, 1137]
Ressource Minimization, Environment Surveillance/Protection	[1138]
Combinatorial Optimization	[1139, 1140, 1141]
Geometry and Physics	[1142, 1143]
Optics and Image Processing	[209, 210, 1144, 204, 205, 1145]

¹ http://en.wikipedia.org/wiki/Evolution_strategy [accessed 2007-07-03], http://www.scholarpedia.org/article/Evolution_Strategies [accessed 2007-07-03]

5.2.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on Evolution Strategy are:

EUROGEN: Evolutionary Methods for Design Optimization and Control with Applications to Industrial Problems
see Section 2.2.2 on page 78

5.2.3 Books

Some books about (or including significant information about) Evolution Strategy are:

Schwefel [440]: *Evolution and Optimum Seeking: The Sixth Generation*
 Rechenberg [105]: *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*
 Rechenberg [106]: *Evolutionsstrategie '94*
 Beyer [1127]: *The theory of evolution strategies*
 Schwefel [1146]: *Numerical Optimization of Computer Models*
 Schöneburg, Heinzmann, and Feddersen [761]: *Genetische Algorithmen und Evolutionsstrategien*
 Bäck [409]: *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*

5.3 Populations in Evolution Strategy

Evolution Strategies usually combine truncation selection (as introduced in Section 2.4.2 on page 94) with one of the following population strategies. These strategies listed below have partly been borrowed from German Wikipedia [2] site for Evolution Strategy².

5.3.1 $(1 + 1)$ -ES

The population only consists of a single individual which is reproduced. From the elder and the offspring, the better individual will survive and form the next population. This scheme is very close to hill climbing which will be introduced in Chapter 8 on page 209.

5.3.2 $(\mu + 1)$ -ES

Here, the population contains μ individuals from which one is drawn randomly. This individual is reproduced from the joint set of its offspring and the current population, the least fit individual is removed.

5.3.3 $(\mu + \lambda)$ -ES

Using the reproduction operations, from μ parent individuals $\lambda \geq \mu$ offspring are created. From the joint set of offspring and parents, only the μ fittest ones are kept [1147].

² <http://de.wikipedia.org/wiki/Evolutionsstrategie> [accessed 2007-07-03]

5.3.4 (μ, λ) -ES

In (μ, λ) Evolution Strategies, introduced by Schwefel [1146], again $\lambda \geq \mu$ children are created from μ parents. The parents are subsequently deleted and from the λ offspring individuals, only the μ fittest are retained [1146, 125].

5.3.5 $(\mu/\rho, \lambda)$ -ES

Evolution Strategies named $(\mu/\rho, \lambda)$ are basically (μ, λ) strategies. The additional parameter ρ is added, denoting the number of parent individuals of one offspring. As already said, normally, we only use mutation ($\rho = 1$). If recombination is also used as in other evolutionary algorithms, $\rho = 2$ holds. A special case of $(\mu/\rho, \lambda)$ algorithms is the $(\mu/\mu, \lambda)$ Evolution Strategy [1148].

5.3.6 $(\mu/\rho + \lambda)$ -ES

Analogously to $(\mu/\rho, \lambda)$ -Evolution Strategies, the $(\mu/\rho + \lambda)$ -Evolution Strategies are (μ, λ) approaches where ρ denotes the number of parents of an offspring individual.

5.3.7 $(\mu', \lambda'(\mu, \lambda)^\gamma)$ -ES

Geyer et al. [1149; 1150; 1151] have developed nested Evolution Strategies where λ' offspring are created and isolated for γ generations from a population of the size μ' . In each of the γ generations, λ children are created from which the fittest μ are passed on to the next generation. After the γ generations, the best individuals from each of the γ isolated solution candidates propagated back to the top-level population, i. e., selected. Then, the cycle starts again with λ' new child individuals. This nested Evolution Strategy can be more efficient than the other approaches when applied to complex multimodal fitness environments [106, 1151].

5.4 One-Fifth Rule

The $\frac{1}{5}$ success rule defined by Rechenberg [105] states that the quotient of the number of successful mutations (i. e., those which lead to fitness improvements) to the total number of mutations should be approximately $\frac{1}{5}$. If the quotient is bigger, the σ -values should be increased and with that, the scatter of the mutation. If it is lower, σ should be decreased and thus, the mutations are narrowed down.

5.5 Differential Evolution

5.5.1 Introduction

Differential Evolution³ (DE, DES) is a method for mathematical optimization of multidimensional functions that belongs to the group of evolution strategies [1152, 1153, 1154, 1155, 1156, 245, 1157]. Developed by Storn and Price [1158], the DE technique has been invented in order to solve the Chebyshev polynomial fitting problem. It has proven to be a very reliable optimization strategy for many different tasks where parameters that can be encoded in real vectors.

The essential idea behind Differential Evolution is the way the recombination operator “deRecombination” is defined for creating new solution candidates. The difference $\mathbf{x}_1 - \mathbf{x}_2$

³ http://en.wikipedia.org/wiki/Differential_evolution [accessed 2007-07-03], <http://www.icsi.berkeley.edu/~storn/code.html> [accessed 2007-07-03]

of two vectors $\mathbf{x}_1$ and $\mathbf{x}_2$ in $\mathbb{X}$ is weighted with a weight $w \in \mathbb{R}$ and added to a third vector $\mathbf{x}_3$ in the population.

$$\mathbf{x} = \text{deRecombination}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \Rightarrow \mathbf{x} = \mathbf{x}_3 + w(\mathbf{x}_1 - \mathbf{x}_2) \quad (5.1)$$

Except for determining w , no additional probability distribution has to be used and the Differential Evolution scheme is completely self-organizing. This classical reproduction strategy has been complemented with new ideas like triangle mutation and alternations with weighted directed strategies.

Gao and Wang [1159] emphasize the close similarities between the reproduction operators of Differential Evolution and the search step of the downhill simplex. Thus, it is only logical to combine or to compare the two methods (see Section 11.4 on page 226). Further improvements to the basic Differential Evolution scheme have been contributed, for instance, by Kaelo and Ali. Their DERL and DELB algorithms outperformed [1160, 1161, 1162] standard DE on the test benchmark from Ali et al. [1163].

5.5.2 General Information

Areas Of Application

Some example areas of application of Differential Evolution are:

Application	References
Engineering, Structural Optimization, and Design	[1164, 1165]
Chemistry, Chemical Engineering	[1166, 1167, 1168, 1169]
Scheduling	[1170]
Function Optimization	[1171]
Electrical Engineering and Circuit Design	[1172, 1173]

Journals

Some journals that deal (at least partially) with Differential Evolution are:

Journal of Heuristics (see Section 1.6.3 on page 64)

Books

Some books about (or including significant information about) Differential Evolution are:

Price, Storn, and Lampinen [1152]: *Differential Evolution – A Practical Approach to Global Optimization*

Feoktistov [1153]: *Differential Evolution – In Search of Solutions*

Corne, Dorigo, Glover, Dasgupta, Moscato, Poli, and Price [400]: *New Ideas in Optimisation*

Evolutionary Programming

6.1 Introduction

Different from the other major types of evolutionary algorithms introduced, there exists no clear specification or algorithmic variant for evolutionary programming¹ (EP) to the knowledge of the author. There is though a semantic difference: while single individuals of a species are the biological metaphor for solution candidates in other evolutionary algorithms, in evolutionary programming, a solution candidate is thought of as a species itself. Hence, mutation and selection are the only operators used in EP and recombination is not applied. The selection scheme utilized in evolutionary programming is quite similar to the $(\mu + \lambda)$ method in Evolution Strategies.

Evolutionary programming was first invented by Lawrence Fogel et al. [785], [56] in order to evolve finite state machines as predictors for data streams. His son David Fogel [1174; 1175; 1176] and he together [1177, 1178] are also the major forces developing it further.

Generally, it is hard to distinguish evolutionary programming from EP and Genetic Programming, genetic algorithms, and Evolution Strategy. Although there are semantic differences (as already mentioned), the author thinks that the evolutionary programming approach generally has melted together with these other research areas.

6.2 General Information

6.2.1 Areas Of Application

Some example areas of application of evolutionary programming are:

Application	References
Machine Learning	[1174]
Cellular Automata and Finite State Machines	[785]
Evolving Behaviors, e.g., for Agents or Game Players	[1175, 1176]
Machine Learning	[1178]
Chemistry, Chemical Engineering and Biochemistry	[1179, 1180, 1181]
Electrical Engineering and Circuit Design	[1182, 1183]
Data Mining and Data Analysis	[1184]
Robotics	[1185]

¹ http://en.wikipedia.org/wiki/Evolutionary_programming [accessed 2007-07-03]

6.2.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on evolutionary programming are:

EP: International Conference on Evolutionary Programming

now part of *CEC*, see Section 2.2.2 on page 77

History: 1998: San Diego, California, USA, see [1186]

1997: Indianapolis, Indiana, USA, see [1187]

1996: San Diego, California, USA, see [1188]

1995: San Diego, California, USA, see [1189]

1994: see [1190]

1993: see [1191]

1992: see [1192]

EUROGEN: Evolutionary Methods for Design Optimization and Control with Applications to Industrial Problems

see Section 2.2.2 on page 78

6.2.3 Books

Some books about (or including significant information about) evolutionary programming are:

Fogel, Owens, and Walsh [785]: *Artificial Intelligence through Simulated Evolution*

Fogel [1174]: *System Identification through Simulated Evolution: A Machine Learning Approach to Modeling*

Fogel [1176]: *Blondie24: playing at the edge of AI*

Bäck [409]: *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*

Learning Classifier Systems

7.1 Introduction

In the late 1970s, Holland invented and patented a new class of cognitive systems, called classifier systems (CS) [1067, 1193, 1068]. These systems are a special case of production systems [1194, 1195] and consist of four major parts:

1. a set of interacting productions, called *classifiers*,
2. a performance algorithm that directs the action of the system in the environment,
3. a simple learning algorithm that keeps track on each classifier's success in bringing about rewards, and
4. a more complex algorithm, called the *genetic algorithm*, that modifies the set of classifiers so that variants of good classifiers persist and new, potentially better ones are created in a provably efficient manner [1196].

By time, classifier systems have undergone some name changes. In 1986, reinforcement learning was added to the approach and the name changed to Learning Classifier Systems¹ (LCS) [1, 1197]. Learning Classifier Systems are sometimes subsumed under a machine learning paradigm called evolutionary reinforcement learning (ERL) [1] or Evolutionary Algorithms for Reinforcement Learning (EARLs) [1198].

7.2 General Information

7.2.1 Areas Of Application

Some example areas of application of Learning Classifier Systems are:

Application	References
Data Mining and Data Analysisg	[1199, 1200, 1201]
Grammar Induction	[1202, 1203, 1204]
Medicine	[1205]
Image Processing	[1206, 1207]
Sequence Prediction	[1208]

7.2.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on Learning Classifier Systems are:

¹ http://en.wikipedia.org/wiki/Learning_classifier_system [accessed 2007-07-03]

IWLCS: International Workshop on Learning Classifier Systems

Nowadays often co-located with GECCO (see Section 2.2.2 on page 79).

History: 2007: London, England, see [1209]
 2006: Seattle, WA, USA, see [1210]
 2005: Washington DC, USA, see [1211, 1212]
 2004: Seattle, Washington, USA, see [1213, 1212]
 2003: Chicago, IL, USA, see [1214, 1212]
 2002: Granada, Spain, see [1215]
 2001: San Francisco, CA, USA, see [1216]
 2000: Paris, France, see [1217]
 1999: Orlando, Florida, USA, see [1218]
 1992: Houston, Texas, USA, see [1219]

7.2.3 Books

Some books about (or including significant information about) Learning Classifier Systems are:

Bull [1220]: *Applications Of Learning Classifier Systems*

Bull and Kovacs [1221]: *Foundations of Learning Classifier Systems*

Butz [1222]: *Anticipatory Learning Classifier Systems*

Butz [1223]: *Rule-Based Evolutionary Online Learning Systems: A Principled Approach to LCS Analysis and Design*

Lanzi, Stolzmann, and Wilson [1224]: *Learning Classifier Systems, From Foundations to Applications*

7.3 The Basic Idea of Learning Classifier Systems

Figure 7.1 illustrates the structure of a Michigan-style Learning Classifier System. A classifier system is connected via detectors (*b*) and effectors (*c*) to its environment (*a*). The input in the system (coming from the detectors) is encoded in form of binary messages that are written into a message list (*d*). On this list, simple **if-then** rules (*e*), the so-called classifiers, are applied. The result of a classification is again encoded as a message and written to the message list. These new messages may now trigger other rules or are signals for the effectors [1226]. The payoff of the performed actions is distributed by the credit apportionment system (*f*) to the rules. Additionally, a rule discovery system (*g*) is responsible for finding new rules and adding them to the classifier population [1225].

Classifier systems are special instances of production systems, which were shown to be Turing complete by Post [1093] and Minsky [1227; 1228]. Thus, Learning Classifier Systems are as powerful as any other Turing-equivalent programming language and can be pictured as something like computer programs where the rules play the role of the instructions and the messages are the memory.

In order to describe how rules and messages are structured in a basic classifier systems, we borrow a simple example from Heitkötter and Beasley [1] and modify it so it fits to our needs. We will orient our explanation at the syntax described by Geyer-Schulz [1225]. You should, however, be aware that there are very many different forms of classifier system and take this as an example for *how it could be done* rather than as *the way it is to be done*.

7.3.1 Messages

So let us imagine that we want to find a classifier system that is able to control the behavior of a frog. Our frog likes to eat nutritious flies. Therefore, it can detect small, flying objects




Fig. 7.1: The structure of a Michigan style Learning Classifier System according to Geyer-Schulz [1225].

and eat them if they are right in front of it. The frog also has a sense of direction and can distinguish between objects that are in front, to the left, or to the right of it and may also turn into any of these directions. The frog can furthermore distinguish objects with stripes from those without. Flying objects with stripes are most likely bees or wasps which should, preferably, not be eaten. The frog can also sense large, looming objects far above, birds in human-lingo, which should be avoided by jumping away quickly. We can compile an according behavior into the form of simple **if-then** rules which are listed in Table 7.1.

No.	if-part	then-part
1	small, flying object with no stripes to the left	send <i>a</i>
2	small, flying object with no stripes to the right	send <i>b</i>
3	small, flying object with no stripes to the front	send <i>c</i>
4	large, looming object	send <i>d</i>
5	<i>a</i> and not <i>d</i>	turn left
6	<i>b</i> and not <i>d</i>	turn right
7	<i>c</i> and not <i>d</i>	eat
8	<i>d</i>	move away rapidly

Table 7.1: if-then rules for frogs




Fig. 7.2: One possible encoding of messages for a frog classifier system

In Figure 7.2, we demonstrate how the messages in a classifier system that drives such a frog can be encoded. Here, input information as well as action commands are compiled in one message type. Also, three bits are assigned for encoding the internal messages a to d . Two bits would not suffice, since 00 occurs in all “original” input messages. At the beginning of a classification process, the input messages are written to the message list. They contain information only at the positions reserved for detections and have zeros in the bits for memory for actions. The classifiers transform them to internal messages which normally have only the bits marked as “memory” set. These messages are finally transformed to output messages by setting some action bits. In our frog-system, a message is in total $k = 12$ bits long, i. e., $|m| = 12 \forall \text{message } m$.

7.3.2 Conditions

Rules in classifier systems consist of a condition part and an action part. The conditions have the same length k as the messages. Instead of being binary encoded strings, a ternary system consisting of the symbols 0, 1, and * is used. In a condition,

- a 0 means that the corresponding bit in the message must be 0,
- a 1 means that the corresponding bit in the message must be 1, and
- a * means *don't care*, i. e., the corresponding bit in the message may be 0 as well as 1 for the condition to match.

Definition 7.1 (match). A message m matches to a condition c if $\text{match}(m, c)$ evaluates to **true**.

$$\text{match}(m, c) = \forall 0 \leq i < |m| \Rightarrow m[i] = c[i] \vee c[i] = * \quad (7.1)$$

The conditional part of a rule may consist of multiple conditions which are implicitly concatenated with logical and ($\wedge$). A classifier is satisfied if all its conditions are satisfied by at least one message in the current message list. This means that each of the conditions of a classifier may match to different messages. We can precede each single condition c with an additional ternary digit which defines if it should be negated or not: * stands for the negation

$\bar{c}$ and 0 as well as 1 denotes c . Here we deviate from one of the syntax described in Geyer-Schulz [1225] because the definition of the “conditionSpecificity” (see Definition 7.2) becomes more beautiful this way. A negated condition evaluates to **true** if no message exists that matches it. By combining **and** and **not**, we get **nands** with which we can build all other logic operations and, hence, whole computers [1229]. Algorithm 7.1 illustrates how the condition part C is matched against the message list M . If the matching is successful, it returns the list S of messages that satisfied the conditions. Otherwise, the output will be the empty list $()$.

Algorithm 7.1: $S \leftarrow \text{matchesConditions}(M, C)$

Input: M : the message list
Input: C : the condition part of a classifier
Input: [implicit] k : the length of the messages $m \in M$ and the single conditions $c \in C$
Input: [implicit] *havePrefix*: **true** if and only if the single conditions have a prefix which determines whether or not they are negated, **false** if no such prefixes are used
Data: i : a counter variable
Data: c : a condition
Data: *neg*: should the condition be negated?
Data: m : a single message from M
Data: b : a Boolean variable
Output: S : the messages that match the condition part C , or $()$ if none such message exists

```

1 begin
2    $S \leftarrow ()$ 
3    $b \leftarrow \text{true}$ 
4    $i \leftarrow 0$ 
5   while  $(i < \text{len}(C)) \wedge b$  do
6     else if havePrefix then
7        $neg \leftarrow (C[i] = *)$ 
8        $i \leftarrow i + 1$ 
9      $neg \leftarrow \text{false}$ 
10     $c \leftarrow \text{subList}(C, i, k)$ 
11     $i \leftarrow i + k$ 
12    else if  $\exists m \in M : \text{match}(m, c)$  then
13       $b \leftarrow \overline{neg}$ 
14      if  $b$  then  $S \leftarrow \text{addListItem}(S, m)$ 
15     $b \leftarrow neg$ 
16    if  $b$  then  $S \leftarrow \text{addListItem}(S, \text{createList}(k, 0))$ 
17  if  $b$  then return  $S$ 
18  else return  $()$ 
19 end
```

Definition 7.2 (Condition Specificity). The condition specificity $\text{conditionSpecificity}(x)$ of a classifier x is the number of non- $*$ symbols in its condition part $C(x)$.

$$\text{conditionSpecificity}(x) = |\{ \forall i : C(x)[i] \neq * \}| \quad (7.2)$$

A rule x_1 with a higher condition specificity is more specific than another rule x_2 with a lower condition specificity. On the other hand, a rule x_2 with $\text{conditionSpecificity}(x_1) > \text{conditionSpecificity}(x_2)$ is more general than the rule x_1 . We can use this information if two rules match to one message, and only one should be allowed to post a message. Preferring the more specific rule in such situations leads to *default hierarchies* [1230, 1231, 1232, 1233] which allows general classifications to “delegate” special cases to specialized classifiers. Even more specialized classifiers can then represent exceptions to these refined rules.

7.3.3 Actions

The action part of a rule has normally exactly the same length as a message. It can either be binary or ternary encoded. In the first case, the action part of a rule is simple copied to the message list if the classifier is satisfied. In the latter case, some sort of merging needs to be performed. Here,

1. a 0 in the action part will lead to a 0 in the corresponding message bit,
2. a 1 in the action part will lead to a 1 in the corresponding message bit,
3. and for a * in the action part, we copy the corresponding bit from the (first) message that matched the classifier's condition to the newly created message.

Definition 7.3 (*mergeAction*). The function “mergeAction” computes a new message n as product of an action a . If the alphabet the action is based on is ternary and may contain *-symbols, mergeAction needs access to the message m which has satisfied the first condition of the classifier to which a belongs. If the classifier contains negation symbols and the first condition was negated, m is assumed to be a string of zeros ($m = \text{createList}(\text{len}(a), 0)$). Notice that we do not explicitly distinguish between binary and ternary encoding in mergeAction, since * cannot occur in actions based on a binary alphabet and Equation 7.3 stays valid.

$$\begin{aligned}
 n = \text{mergeAction}(a, m) \Leftrightarrow & (\text{len}(n) = \text{len}(a)) \wedge \\
 & (n[i] = a[i] \ \forall i \in 0..\text{len}(a) - 1 : a[i] \neq *) \wedge \\
 & (n[i] = m[i] \ \forall i \in 0..\text{len}(a) - 1 : a[i] = *) \quad (7.3)
 \end{aligned}$$

7.3.4 Classifiers

So we know that a rule x consist of a condition part $C(x)$ and an action part $a(x)$. C is list of $r \in \mathbb{N}$ conditions c_i , and we distinguish between representations with ($C = (n_1, c_1, n_2, c_2, \dots, n_r, c_r)$) and without negation symbol ($C = (c_1, c_2, \dots, c_r)$). Let us now go back to our frog example. Based on the encoding scheme defined in Figure 7.2, we can translate Table 7.1 into a set of classifiers. In the example, we use two conditions c_1 and c_2 with the negation symbols n_1 and n_2 , i. e., $r = 2$. Table 7.2 contains the result of this

No.	n_1	c_1	n_2	c_2	a
1	0	0 0 0 01 0 *** ** *	0	* * * * * * * * *	0 0 00 0 001 00 0 0
2	0	0 0 0 11 0 *** ** *	0	* * * * * * * * *	0 0 00 0 010 00 0 0
3	0	0 0 0 10 0 *** ** *	0	* * * * * * * * *	0 0 00 0 011 00 0 0
4	0	1 1 * * * * * * *	0	* * * * * * * * *	0 0 00 0 100 00 0 0
5	0	* * * * * 001 ** *	*	* * * * * 100 ** *	0 0 00 0 000 01 0 0
6	0	* * * * * 010 ** *	*	* * * * * 100 ** *	0 0 00 0 000 10 0 0
7	0	* * * * * 011 ** *	*	* * * * * 100 ** *	0 0 00 0 000 00 0 1
8	0	* * * * * 100 ** *	0	* * * * * * * * *	0 0 00 0 000 00 1 0

Table 7.2: if-then rules for frogs in encoded form

encoding. We can apply this classifier to a situation in the life of our frog where it detects

1. a fly to its left,
2. a bee to its right, and
3. a stork left in the air.

How will it react? The input sensors will generate three messages and insert them into the message list $M_1 = (m_1, m_2, m_3)$:

1. $m_1 = (000100000000)$ for the fly,

2. $m_2 = (001110000000)$ for the bee, and
3. $m_3 = (110100000000)$ for the stork.

The first message triggers rule 1 and the third message triggers rule 4 whereas no condition fits to the second message. As a result, the new message list M_2 contains two messages, m_4 and m_5 , produced by the corresponding actions.

1. $m_4 = (000000010000)$ from rule 1 and
2. $m_5 = (000001000000)$ from rule 4.

m_4 could trigger rule 5 but is inhibited by the negated second condition c_2 because of message m_5 . m_5 matches to classifier 8 which finally produces message $m_6 = (000000000010)$ which forces the frog to jump away. No further classifiers become satisfied with the new message list $M_3 = (m_6)$ and the classification process is terminated.

7.3.5 Non-Learning Classifier Systems

So far, we have described a non-learning classifier system. Algorithm 7.2 defines the behavior of such a system which we also could observe in the example. It still lacks the credit apportionment and the rule discovery systems (see (f) and (g) in Figure 7.1). A non-learning classifier is able to operate correctly on a fixed set of situations. It is sufficient for all applications where we are able to determine this set beforehand and no further adaptation is required. If this is the case, we can use genetic algorithms to evolve the classifier systems offline, for instance.

Algorithm 7.2: nonLearningClassifierSystem(P)

Input: P : the list of rules x_i that determine the behavior of the classifier system
Input: [implicit] readDetectors: a function which creates a new message list containing only the input messages from the detectors
Input: [implicit] sendEffectors: a function which translates all messages concerning effectors to signals for the output interface
Input: [implicit] $t_{max} \in \mathbb{N}$: the maximum number of iterations for the internal loop, avoids endless loops
Data: t : a counter the internal loop
Data: M, N, S : the message lists
Data: x : a single classifier

```

1 begin
2   while true do
3      $M \leftarrow \text{readDetectors}()$ 
4      $t \leftarrow 0$ 
5     repeat
6        $N \leftarrow ()$ 
7       foreach  $x \in P$  do
8          $S \leftarrow \text{matchesConditions}(M, C(x))$ 
9         if  $\text{len}(S) > 0$  then
10           $N \leftarrow \text{addListItem}(N, \text{mergeAction}(a(x), S_{[0]}))$ 
11        $M \leftarrow N$ 
12        $t \leftarrow t + 1$ 
13     until  $(\text{len}(M) = 0) \vee (t > t_{max})$ 
14     if  $\text{len}(M) > 0$  then sendEffectors( $M$ )
15 end
```

Algorithm 7.2 illustrates how a classifier system works. No optimization or approximation of a solution is done; this is a complete control system in action. Therefore we do not need a termination criterion but run an infinite loop.

7.3.6 Learning Classifier Systems

In order to convert this non-learning classifier system to Learning Classifier System as proposed by Holland [1234] and sketched in Algorithm 7.3, we have to add the aforementioned missing components. Heitkötter and Beasley [1] suggest two ways for doing so:

1. Currently, the activation of a classifier x results solely from the message-matching process. If a message matches the condition(s) $C(x)$, the classifier may perform its action $a(x)$. We can change this mechanism by making it also dependent on an additional parameter $v(x)$ – a strength value, which can be modified as a result of experience, i. e., by reinforcement from the environment. Therefore, we have to solve the *credit assignment problem* first defined by Minsky [1235; 1236], since chains of multiple classifiers can cause a certain action.
2. Furthermore (or instead), we may also modify the set of classifiers P by adding, removing, or combining condition/action parts of existing classifiers.

A Learning Classifier System hence is a control system which is able to learn while actually running and performing its work. Usually, a training phase will precede any actual deployment. Afterwards, the learning may even be deactivated, which turns the LCS into an ordinary classifier system or the learning rate is decreased.

7.3.7 The Bucket Brigade Algorithm

The Bucket Brigade Algorithm has been developed by Holland [1237; 1234] as one method of solving the credit assignment problem in Learning Classifier Systems. Research work concerning this approach and its possible extensions has been conducted by Westerdale [1238; 1239; 1240], Antonisse [1241], Huang [1242], Riolo [1243; 1231], Dorigo [1244], Spiessens [1245], Wilson [1246], Holland and Burks [1068], and Hewahi and Bharadwaj [1247] and has neatly been summarized by Hewahi [1248], [1249]. In the following, we will outline this approach with the notation of de Boer [1226].

The Bucket Brigade Algorithm selects the classifiers from the match set X that are allowed to post a message (i. e., becoming member in the activated set U) by an auction. Therefore, each matching classifier x places a bid $B(x)$ which is the product of a linear function ϑ of the condition specificity of x , a constant $0 < \beta \leq 1$ that determines the fraction of the strength of x should be used and its strength $v(x)$ itself. In practical applications, values like $\frac{1}{8}$ or $\frac{1}{16}$ are often chosen for β .

$$B(x) = \vartheta(x) * \beta * v(x) + \text{random}_n(0, \sigma^2) \quad (7.4)$$

Sometimes, a normal distributed random number is added to each bid in order to make the decisions of the system less deterministic, as done in Equation 7.4.

The condition specificity is included in the bid calculation because it gives a higher value to rules with fewer *-symbols in their conditions. These rules match to fewer messages and can be considered more relevant in the cases they do match. For ϑ , the quotient of the number non-*-symbols and the condition length plus some constant $0 < \alpha$ determining the importance of the specificity of the classifier is often used [1226].

$$\vartheta(x) = \frac{\text{conditionSpecificity}(x)}{\text{len}(C(x))} + \alpha \quad (7.5)$$

The bucket brigade version of the selectMatchingClassifiers-function introduced in Algorithm 7.3 then picks the k classifiers with the highest bids and allows them to write

Algorithm 7.3: learningClassifierSystem()

Input: P : the list of rules x_i that determine the behavior of the classifier system
Input: [implicit] generateClassifiers: a function which creates randomly a population P of classifiers
Input: [implicit] readDetectors: a function which creates a new message list containing only the input messages from the detectors
Input: [implicit] sendEffectors: a function which translates all messages concerning effectors to signals for the output interface
Input: [implicit] selectMatchingClassifiers: a function that determines at most k classifiers from the matching set that are allowed to trigger their actions
Input: [implicit] generationCriterion: a criterion that becomes **true** if new classifiers should be created
Input: [implicit] updateRules: a function that finds new rules and deletes old ones
Input: [implicit] $t_{max} \in \mathbb{N}$: the maximum number of iterations for the internal loop, avoids endless loops
Data: t, i : counter variables
Data: M, N, S : the message lists
Data: X : a list of tuples containing classifiers and the (first) messages that satisfied their conditions
Data: v : the strength values
Data: x : a single classifier

```

1 begin
2    $P \leftarrow \text{generateClassifiers}(s)$ 
3   foreach  $x \in P$  do  $v(x) \leftarrow 1$ 
4   while true do
5      $M \leftarrow \text{readDetectors}()$ 
6      $t \leftarrow 0$ 
7     repeat
8        $X \leftarrow ()$ 
9       foreach  $x \in P$  do
10         $S \leftarrow \text{matchesConditions}(M, C(x))$ 
11        if  $\text{len}(S) > 0$  then
12           $X \leftarrow \text{addListItem}(X, (x, S_{[0]}))$ 
13       $N \leftarrow ()$ 
14      if  $\text{len}(X) > 0$  then
15         $(X, v) \leftarrow \text{selectMatchingClassifiers}(X, v)$ 
16        for  $i \leftarrow 0$  up to  $\text{len}(X) - 1$  do
17           $x \leftarrow X_{[i,0]}$ 
18           $N \leftarrow \text{addListItem}(N, \text{mergeAction}(a(x), X_{[i,1]}))$ 
19       $M \leftarrow N$ 
20       $t \leftarrow t + 1$ 
21    until  $(\text{len}(M) = 0) \vee (t > t_{max})$ 
22    if  $\text{len}(M) > 0$  then
23      sendEffectors( $M$ )
24      // distribute Payoffs
25      if generationCriterion() then  $P \leftarrow \text{updateRules}(P, v)$ 
25 end
  
```

their messages into the new message list. They are charged with the payment part $P(x)$ of their bids. The payment does not contain the condition specificity-dependent part and also not the possible random addend. It is added as reward $R(y)$ to the strength of classifier y that wrote the message which allowed them to become active. In the case that this was an input message, it is simply thrown away. The payment of classifiers that are not activated is null.

$$P(x) = \beta * v(x) \quad (7.6)$$

In some Learning Classifier Systems, a life-tax $T(x)$ is collected from all classifiers in each cycle. It is computed as a small fraction τ of their strength.

$$T(x) = \tau * v(x) \quad (7.7)$$

Those classifiers that successfully triggered an action of the effectors receive a reward $R(x)$ from the environment which is added to their strength. Together with the payment method, all rules that are involved in a successful action receive some of the reward which is handed down stepwise – similar to how water is transported by a bucket brigade. For all classifiers that do not produce output to the effectors and also do not receive payment from other classifier they have triggered, this reward is null.

In total, the new strength $v_{t+1}(x)$ of a classifier x is composed of its old strength, its payment $P(x)$, the life-tax $T(x)$, and the reward $R(x)$.

$$v_{t+1}(x) = v_t(x) - P(x) - T(x) + R(x) \quad (7.8)$$

Instead of the Bucket Brigade Algorithm, it is also possible to use Q-Learning in Learning Classifier Systems, as shown by Wilson [1250]. Dorigo and Bersini [1251] have shown that both concepts are roughly equivalent [1].

7.3.8 Applying the Genetic Algorithm

With the credit assignment alone, no new rules can be discovered – only the initial, randomly create rule set P is rated. At some certain points in time, a genetic algorithm (see Chapter 3 on page 113) replaces old rules by new ones. In Learning Classifier Systems we apply steady-state genetic algorithms which are discussed in Section 2.1.2 on page 73. They will retain most of the classifier population and only replace the weakest rules. Therefore, the strength $v(x)$ of a rule x is directly used as its fitness and is subject to maximization.

For mutation and crossover, the well known reproduction operations for fixed-length string chromosomes discussed in Section 3.4 on page 119 are employed.

7.4 Families of Learning Classifier Systems

The exact definition of Learning Classifier Systems still seems contentious [1252, 1253, 1197, 1254] and there exist many different implementations. There are, for example, versions without message list where the action part of the rules does not encode messages but direct output signals. Also the importance of the role of genetic algorithms in conjunction with the reinforcement learning component is not quite clear. There are scientists who emphasize more the role of the learning components [1255] and others who tend to grant the genetic algorithms a higher weight [1256, 1067]. The families of Learning Classifier Systems have been listed and discussed by Brownlee [1257] elaborately. Here we will just summarize their differences in short. De Jong [1258; 1259] and Grefenstette [1260] divide Learning Classifier Systems into two main types, depending on how the genetic algorithm acts: The *Pitt approach* originated at the University of Pittsburgh with the LS-1 system developed by Smith [1070]. It was then developed further and applied by Spears and De Jong [1069], De Jong and Spears [1261], De Jong et al. [1262], Bacardit i Peñarroya [1200; 1263], and Bacardit

i Peñarroya and Krasnogor [1264]. Pittsburgh-style Learning Classifier Systems work on a population of separate classifier systems, which are combined and reproduced by the genetic algorithm.

The original idea of Holland and Reitman [1067] were *Michigan-style* LCSs, where the whole population itself is considered as classifier system. They focus on selecting the best rules in this rule set [1265, 1226, 1266].

Wilson [1250; 1267] developed two subtypes of Michigan-style LCS:

1. In ZCS systems, there is no message list use fitness sharing [1250, 1268, 1269, 1270] for a Q-learning-like reinforcement learning approach called QBB.
2. ZCS have later been somewhat superseded by XCS systems in which the Bucket Brigade Algorithm has fully been replaced by Q-learning. Furthermore, the credit assignment is based on the accuracy (usefulness) of the classifiers. The genetic algorithm is applied to sub-populations containing only classifiers which apply to the same situations. [1267, 1271, 1272, 1273, 1274]

Hill Climbing

*“Up next we got CaTash with that West coast rhymin’
Bombin niggaz with the style that’s dangerous as mountain climbin’
Background: AaaaaaAaahhh...”*

— Tash [1275]

8.1 Introduction

Hill climbing¹ (HC) [1276] is a very old and simple search and optimization algorithm for single objective functions f . In principle, hill climbing algorithms perform a loop in which the currently known best solution individual p^* is used to produce one offspring p_{new} . If this new individual is better than its parent, it replaces it. Then, the cycle starts all over again. In this sense, it is similar to an evolutionary algorithm with a population size p_{sof} of 1. Although the search space $\mathbb{G}$ and the problem space $\mathbb{X}$ are most often the same in hill climbing, we distinguish them in Algorithm 8.1 for the sake of generality. Hill climbing furthermore normally uses a parameterless search operation to create the first solution candidate and, from there on, unary operations to produce the offspring. Without loss of generality, we will thus make use of the reproduction operations from evolutionary algorithms defined in Section 2.5 on page 108, i. e., set $Op = \{\text{create}, \text{mutate}\}$.

The major problem of hill climbing is premature convergence, i. e., it gets easily stuck on a local optimum. It *always* uses the best known solution candidate x^* to find new points in the problem space $\mathbb{X}$. Hill climbing utilizes a unary reproduction operation similar to mutation in evolutionary algorithms. It should be noted that hill climbing can be implemented in a deterministic manner if the neighbor sets in search space $\mathbb{G}$, which here most often equals the problem space $\mathbb{X}$, are always finite and can be iterated over.

¹ http://en.wikipedia.org/wiki/Hill_climbing [accessed 2007-07-03]

Algorithm 8.1: $x^* \leftarrow \text{hillClimber}(f)$ **Input:** f : the objective function subject to minimization**Data:** p_{new} : the new element created**Data:** p^* : the (currently) best individual**Output:** x^* : the best element found

```

1 begin
2    $p^*.g \leftarrow \text{create}()$ 
   // Implicitly:  $p^*.x \leftarrow \text{gpm}(p^*.g)$ 
3   while  $\text{terminationCriterion}()$  do
4      $p_{\text{new}}.g \leftarrow \text{mutate}(p^*.g)$ 
     // Implicitly:  $p_{\text{new}}.x \leftarrow \text{gpm}(p_{\text{new}}.g)$ 
5     if  $f(p_{\text{new}}.x) < f(p^*.x)$  then  $p^* \leftarrow p_{\text{new}}$ 
6   return  $p^*.x$ 
7 end

```

8.2 General Information

8.2.1 Areas Of Application

Some example areas of application of hill climbing are:

Application	References
Networking and Communication	[1277, 454]
Robotics	[1278]
Data Mining and Data Analysis	[1279]
Evolving Behaviors, e.g., for Agents or Game Players	[1280]
Combinatorial Optimization	[1281, 1282]

8.3 Multi-Objective Hill Climbing

As illustrated in Algorithm 8.2 on the next page, we can easily extend hill climbing algorithms with a support for multi-objective optimization by using some of the methods of evolutionary algorithms. This extended approach will then return a set X^* of the best solutions found instead of a single individual x^* as done in Algorithm 8.1. The set of currently known best individuals Arc may contain more than one element. Therefore, we employ a selection scheme in order to determine which of these individuals should be used as parent for the next offspring in the multi-objective hill climbing algorithm. The selection algorithm applied must not solely rely on the prevalence comparison, since no element in Arc prevails any other. Thus, we also copy the idea of fitness assignment from evolutionary algorithms. For maintaining the optimal set, we apply the updating and pruning methods defined in Chapter 18 on page 263.

Algorithm 8.2: $x^* \leftarrow \text{hillClimberMO}(\text{cmp}_F, a)$

Input: cmp_F : the prevalence comparator
Input: as : the maximum archive size
Data: p_{new} : the new individual generated
Data: Arc : the set of best individuals known
Output: X^* : the set of the best elements found

```

1 begin
2    $Arc \leftarrow ()$ 
3    $p_{new}.x \leftarrow \text{create}()$ 
4   // Implicitly:  $p_{new}.x \leftarrow \text{gpm}(p_{new}.g)$ 
5   while terminationCriterion() do
6      $Arc \leftarrow \text{updateOptimalSet}(Arc, p_{new})$ 
7      $Arc \leftarrow \text{pruneOptimalSet}(Arc, as)$ 
8      $v \leftarrow \text{assignFitness}(Arc, \text{cmp}_F)$ 
9      $p_{new} \leftarrow \text{select}(Arc, v, 1)[0]$ 
10     $p_{new}.g \leftarrow \text{mutate}(p_{new}.g)$ 
11    // Implicitly:  $p_{new}.x \leftarrow \text{gpm}(p_{new}.g)$ 
12  return extractPhenotypes( $Arc$ )
13 end
  
```

8.4 Problems in Hill Climbing

Both versions of the algorithm are still very likely to get stuck on local optima. They will only follow a path of solution candidates if it is monotonously² improving the objective function(s). Hill climbing in this form is a local search rather than global optimization algorithm. By making a few slight modifications to the algorithm however, it can become a valuable global optimization technique:

- A tabu-list which stores the elements recently evaluated can be added. By preventing the algorithm from visiting them again, a better exploration of the problem space $\mathbb{X}$ can be enforced. This technique is used in Tabu Search which is discussed in Chapter 12 on page 229.
- Another way of preventing premature convergence is to not always transcend to the better solution candidate in each step. Simulated Annealing introduces a heuristic based on the physical model the cooling down molten metal to decide whether a superior offspring should replace its parent or not. This approach is described in Chapter 10 on page 217.
- The Dynamic Hill Climbing approach by Yuret and de la Maza [1283] uses the last two visited points to compute unit vectors. With this technique, the directions are adjusted according to the structure of the problem space and a new coordinate frame is created which points more likely into the right direction.
- Randomly restarting the search after so-and-so many steps is a crude but efficient method to explore wide ranges of the problem space with hill climbing. You can find it outlined in Section 8.5.
- Using a reproduction scheme that not necessarily generates solution candidates directly neighboring x^* , as done in Random Optimization, an optimization approach defined in Chapter 9 on page 213, may prove even more efficient.

² <http://en.wikipedia.org/wiki/Monotonicity> [accessed 2007-07-03]

8.5 Hill Climbing with Random Restarts

Hill climbing with random restarts is also called *Stochastic Hill Climbing* (SH) or *Stochastic gradient descent*³ [1284, 1285]. We have mentioned it as a measure for preventing premature convergence; here we want to take a deeper look on this approach.

Let us further combine it directly with the multi-objective hill climbing approach defined in Algorithm 8.2. The new algorithm incorporates two archives for optimal solutions: Arc_1 , the overall optimal set, and Arc_2 the set of the best individuals in the current run. We additionally define the criterion `shouldRestart()` which is evaluated in every iteration and determines whether or not the algorithm should be restarted. `shouldRestart()` therefore could, for example, count the iterations performed or check if any improvement was produced in the last ten iterations. After each single run, Arc_2 is incorporated into Arc_1 , from which we extract and return the problem space elements at the end of the hill climbing process, as defined in Algorithm 8.3.

Algorithm 8.3: $X^* \leftarrow \text{hillClimberMO-RR}(\text{cmp}_F, a)$

Input: cmp_F : the prevalence comparator
Input: a : the maximum archive size
Data: p_{new} : the new individual generated
Data: Arc_1, Arc_2 : the sets of best individuals known
Output: X^* : the set of the best elements found

```

1 begin
2    $Arc_1 \leftarrow ()$ 
   // Implicitly:  $p_{\text{new}}.x \leftarrow \text{gpm}(p_{\text{new}}.g)$ 
3   while terminationCriterion() do
4      $Arc_2 \leftarrow ()$ 
5      $p_{\text{new}}.x \leftarrow \text{create}()$ 
6     while terminationCriterion()  $\vee$  shouldRestart() do
7        $Arc_2 \leftarrow \text{updateOptimalSet}(Arc_2, p_{\text{new}})$ 
8        $Arc_2 \leftarrow \text{pruneOptimalSet}(Arc_2, a)$ 
9        $v \leftarrow \text{assignFitness}(Arc_2, \text{cmp}_F)$ 
10       $p_{\text{new}} \leftarrow \text{select}(Arc_2, v, 1)[0]$ 
11       $p_{\text{new}}.g \leftarrow \text{mutate}(p_{\text{new}}.g)$ 
       // Implicitly:  $p_{\text{new}}.x \leftarrow \text{gpm}(p_{\text{new}}.g)$ 
12       $Arc_1 \leftarrow \text{updateOptimalSetN}(Arc_1, Arc_2)$ 
13       $Arc_1 \leftarrow \text{pruneOptimalSet}(Arc_1, a)$ 
14   return extractPhenotypes( $Arc_1$ )
15 end
```

³ http://en.wikipedia.org/wiki/Stochastic_gradient_descent [accessed 2007-07-03]

Random Optimization

9.1 Introduction

The Random Optimization¹ method for single-objective, numerical problems, i. e., $\mathbb{G} = \mathbb{R}^n$ and $|F| = 1$, was first proposed by Rastrigin [1286] in the early 1960s. It was studied thoroughly by Gurin and Rastrigin [223], Schumer [1287] and further improved by Schumer and Steiglitz [1288] [1289]. A different Random Optimization approach has been introduced by Matyas [1290; 1291] around the same time. Matyas gave theorems about the convergence properties of his approach for unimodal optimization. Baba [1292] then showed theoretically that the global optimum of an optimization problem can even be found if the objective function is multimodal.

There are, however, three important differences between the two approaches:

1. In traditional hill climbing, the new solution candidates are created from a good individual are always very close neighbors of it. In Random Optimization, this is not *necessary* but only *probably*.
2. In Random Optimization, the unary search operation explicitly uses random numbers whereas the unary search operations of hill climbing may be deterministic or randomized.
3. In Random Optimization, the search space $\mathbb{G}$ is always the $\mathbb{R}^n$, the space of n -dimensional real vectors.
4. In Random Optimization, we explicitly distinguish between objective functions $f \in F$ and constraints $c \in C$.

Random Optimization introduces a new search operation “roReproduce” specialized for the numerical search space similar to mutation in Evolution Strategies. This operation is constructed in a way that all points in the search space $\mathbb{G} = \mathbb{R}^n$ can be reached in one step when starting out from every other point. In other words, the operator “roReproduce” is *complete* in the sense of Definition 1.27.

$$\text{roReproduce}(g) = g + \mathbf{r} : \mathbf{r} = \begin{pmatrix} \text{random}_n(\mu_1, \sigma_1^2) & (\sim N(\mu_1, \sigma_1^2)) \\ \text{random}_n(\mu_2, \sigma_2^2) & (\sim N(\mu_2, \sigma_2^2)) \\ \vdots & \vdots \\ \text{random}_n(\mu_n, \sigma_n^2) & (\sim N(\mu_n, \sigma_n^2)) \end{pmatrix} \quad (9.1)$$

Equation 9.1 illustrates one approach to realize such a complete search operation. To the (genotype) of the best solution candidate discovered so far, we add a vector $\mathbf{r} \in \mathbb{R}^n$. Each component $\mathbf{r}[i]$ of this vector is normally distributed around a value μ_i . Hence, the probability density function underlying the components of this vector is greater than zero for all real numbers. The μ_i are the expected values and the σ_i the standard deviations of the normal distributions, as introduced in Section 27.4.2 on page 437. The μ_i define a general direction for the search, i. e., if $\mu_i > 0$, $\text{random}_n(\mu_i, \sigma_i^2)$ will likely also be greater zero and

¹ http://en.wikipedia.org/wiki/Random_optimization [accessed 2007-07-03]

for $\mu_i < 0$, it will probably be smaller than zero, too (given that $|\sigma_i| \ll |\text{expectedValue}M_i|$). The σ_i can be imagined as the range in which the random numbers are distributed around the μ_i and denote a step width of the random numbers. If we choose the absolute values of both, μ_i and σ_i , very small, we can exploit a local optimum whereas larger values lead to a rougher exploration of search space. If the μ s are set to 0, the probability distribution of random numbers will be “centered” around the genotype it is applied to. Since the normal distribution is generally unbounded, it is possible that the random elements in $\mathbf{r}$ can become very large, even for small $\sigma_i \approx 0$. Therefore, local optima can be left again even with bad settings of μ and σ .

Equation 9.1 is only one way to realize the completeness of the “roReproduce”-operation. Instead of the normal distribution, any other probability distribution with $f_X(y) > 0 \forall y \in \mathbb{R}$ would do. Good properties can, for instance, be attributed to the bell-shaped distribution used by Worakul et al. [1293; 1294] and discussed in Section 27.7.3 on page 464 in this book.

In order to respect the idea of constraint optimization in Random Optimization as introduced in [1292], we define a set of constraint functions C . A constraint $c \in C$ is satisfied by a solution candidate $x \in \mathbb{X}$, if $c(x) \leq 0$ holds.

Algorithm 9.1 illustrates how random optimization works, clearly showing connatural traits in comparison with the hill climbing approach Algorithm 8.1 on page 210.

Algorithm 9.1: $x^* \leftarrow \text{randomOptimizer}f$

Input: f : the objective function subject to minimization

Data: p_{new} : the new element created

Data: p^* : the (currently) best individual

Output: x^* : the best element found

```

1 begin
2    $p^*.g \leftarrow \text{create}()$ 
   // Implicitly:  $p^*.x \leftarrow \text{gpm}(p^*.g)$ 
3   while  $\text{terminationCriterion}()$  do
4      $p_{\text{new}}.g \leftarrow \text{roReproduce}(p^*.g)$ 
     // Implicitly:  $p_{\text{new}}.x \leftarrow \text{gpm}(p_{\text{new}}.g)$ 
5     if  $c(p_{\text{new}}.x) \leq 0 \forall c \in C$  then
6       if  $f(p_{\text{new}}.x) < f(p^*.x)$  then  $p^* \leftarrow p_{\text{new}}$ 
7   return  $p^*.x$ 
8 end
```

Setting the values of μ and σ adaptively can lead to large improvements in convergence speed. The Heuristic Random Optimization (HRO) algorithm introduced by Li and Rhinehart [1295] and its successor method Random Optimization II developed by Chandran and Rhinehart [1296] for example update them by utilizing gradient information or reinforcement learning.

9.2 General Information

9.2.1 Areas Of Application

Some example areas of application of (heuristic) Random Optimization are:

Application

References

Medicine	[1293, 1294]
Biology and Medicine	[1297]
Machine Learning	[1298]
Function Optimization	[1295, 1299]

Simulated Annealing

*“Now if you were to alter masks every time fame circus approaches
Do you really think your maker wouldn’t notice?”*

— Aesop Rock [779]

10.1 Introduction

In 1953, Metropolis et al. [1300] developed a Monte Carlo method for “calculating the properties of any substance which may be considered as composed of interacting individual molecules”. With this so-called “Metropolis” procedure stemming from statistical mechanics, the manner in which molten metal crystallizes and reaches equilibria in the process of annealing can, for instance, be simulated. This inspired Kirkpatrick et al. [1301] to develop the Simulated Annealing¹ (SA) global optimization algorithm in the early 1980s and to apply it to various combinatorial optimization problems. Independently, Černý [1302] applied a similar approach to the travelling salesman problem. Simulated Annealing is an optimization method that can be applied to arbitrary search and problem spaces. It only needs one parameterless search operation for creating the starting point of the search and a unary operation for the subsequent search steps.

Annealing steel is cooled down slowly in order to keep the system of the melt in a thermodynamic equilibrium which will increase the size of its crystals and reduce their defects. As cooling proceeds, the atoms of the steel become more ordered. At some point of time, the system reaches its ground state and the metal has solidified, i.e., “freezes”. Beyond this point, no changes in the arrangement of the atoms will occur. For all possible liquids, this will happen at absolute zero, the thermodynamic temperature $T = 0\text{K}$ which is the lowest energy state possible, at the latest. When annealing metal, the initial temperature must not be too low and the cooling must be done sufficiently slowly so as to avoid the system getting stuck in a meta-stable, non-crystalline, state representing a local minimum of energy.

In physics, each set of positions of all atoms of a system pos is weighted by its Boltzmann probability factor $e^{-\frac{E(pos)}{k_B T}}$ where $E(pos)$ is the energy of the configuration pos , T is the temperature measured in Kelvin, and k_B is the Boltzmann’s constant² $k_B = 1.380\,650\,524 \cdot 10^{-23}\text{J/K}$.

The Metropolis procedure was an exact copy of this physical process which could be used to simulate a collection of atoms in thermodynamic equilibrium at a given temperature. A new nearby geometry pos_{i+1} was generated as a random displacement from the current geometry pos_i of an atom in each iteration. The energy of the resulting new geometry is computed and ΔE , the energetic difference between the current and the new geometry, was determined. The probability that this new geometry is accepted, $P(\Delta E)$ is defined in Equation 10.2.

¹ http://en.wikipedia.org/wiki/Simulated_annealing [accessed 2007-07-03]

² http://en.wikipedia.org/wiki/Boltzmann%27s_constant [accessed 2007-07-03]

$$\Delta E = E(pos_{i+1}) - E(pos_i) \quad (10.1)$$

$$P(\Delta E) = \begin{cases} e^{-\frac{\Delta E}{k_B T}} & \text{if } \Delta E > 0 \\ 1 & \text{otherwise} \end{cases} \quad (10.2)$$

Thus, if the new nearby geometry has a lower energy level, the transition is accepted. Otherwise, a uniformly distributed random number $r = \text{random}_u() \in [0, 1)$ is drawn and the step will only be accepted in the simulation if it is less or equal the Boltzmann probability factor, i.e., $r \leq P(\Delta E)$. At high temperatures T , this factor is very close to 1, leading to the acceptance of many uphill steps. As the temperature falls, the proportion of steps accepted which would increase the energy level decreases. Now the system will not escape local regions anymore and (hopefully) comes to a rest in the global minimum at temperature $T = 0\text{K}$.

The abstraction of this method in order to allow arbitrary problem spaces is straightforward – the energy computation $E(pos_i)$ is replaced by an objective function f or even by the result v of a fitness assignment process. Algorithm 10.1 illustrates the basic course of Simulated Annealing. Without loss of generality, we reuse the definitions from evolutionary algorithms for the search operations and set $Op = \{\text{create}, \text{mutate}\}$.

Algorithm 10.1: $x^* \leftarrow \text{simulatedAnnealing}(f)$

Input: f : the objective function to be minimized
Data: p_{new} : the newly generated individual
Data: p_{cur} : the point currently investigated in problem space
Data: p^* : the best individual found so far
Data: T : the temperature of the system which is decreased over time
Data: t : the current time index
Data: ΔE : the energy difference of the x_{new} and x_{cur}
Output: x^* : the best element found

```

1 begin
2    $p_{new}.g \leftarrow \text{create}()$ 
   // Implicitly:  $p_{new}.x \leftarrow \text{gpm}(p_{new}.g)$ 
3    $p_{cur} \leftarrow p_{new}$ 
4    $p^* \leftarrow p_{new}$ 
5    $t \leftarrow 0$ 
6   while  $\text{terminationCriterion}()$  do
7      $\Delta E \leftarrow f(p_{new}.x) - f(p_{cur}.x)$ 
8     if  $\Delta E \leq 0$  then
9        $p_{cur} \leftarrow p_{new}$ 
10      if  $f(p_{cur}.x) < f(p^*.x)$  then  $p^* \leftarrow p_{cur}$ 
11    else
12       $T \leftarrow \text{getTemperature}(t)$ 
13      if  $\text{random}_u() < e^{-\frac{\Delta E}{k_B T}}$  then  $p_{cur} \leftarrow p_{new}$ 
14       $p_{new}.g \leftarrow \text{mutate}(p_{cur}.g)$ 
   // Implicitly:  $p_{new}.x \leftarrow \text{gpm}(p_{new}.g)$ 
15       $t \leftarrow t + 1$ 
16   return  $p^*.x$ 
17 end
```

It has been shown that Simulated Annealing algorithms with appropriate cooling strategies will asymptotically converge to the global optimum. Nolte and Schrader [1303] and Laarhoven and Aarts [1304] provide lists of the most important works showing that Simulated Annealing will converge to the global optimum if $t \rightarrow \infty$ iterations are performed,

including the studies of Hajek [1305]. Nolte and Schrader [1303] further list research providing deterministic, non-infinite boundaries for the asymptotic convergence by Anily and Federgruen [1306], Gidas [1307], Nolte and Schrader [1308], Mitra et al. [1309]. In the same paper, they introduce a significantly lower bound, which, however, still states that Simulated Annealing is probably in an optimal configuration after the number of iterations exceeds the cardinality of the problem space – which is, well, slow [1310]. In other words, it would be faster to enumerate all possible solution candidates in order to find the global optimum with absolute certainty than applying Simulated Annealing. This does not mean that Simulated Annealing is always slow. It only needs that much time if we persist on the optimality. Speeding up the cooling process will result in a faster search, but voids the guaranteed convergence on the other hand. Such speeded-up algorithms are called *Simulated Quenching* (SQ) [1311, 1312, 1313].

10.2 General Information

10.2.1 Areas Of Application

Some example areas of application of Simulated Annealing are:

Application	References
Combinatorial Optimization	[1302, 1301, 1314]
Function Optimization	[1315]
Chemistry, Chemical Engineering	[1316, 1317, 690, 691]
Image Processing	[1318, 1319, 1320, 1314]
Economics and Finance	[1321, 1322]
Electrical Engineering and Circuit Design	[1323, 1314]
Machine Learning	[1324, 1325, 1314, 1326]
Geometry and Physics	[1327, 1314, 1328, 1329]
Networking and Communication	[717]

For more information see also [1304].

10.2.2 Books

Some books about (or including significant information about) Simulated Annealing are:

Laarhoven and Aarts [1304]: <i>Simulated Annealing: Theory and Applications</i>
Badiru [1330]: <i>Handbook of Industrial and Systems Engineering</i>
Davis [765]: <i>Genetic Algorithms and Simulated Annealing</i>

10.3 Temperature Scheduling

The temperature schedule defines how the temperature in Simulated Annealing is decreased. As already mentioned, this has major influence on whether the Simulated Annealing algorithm will succeed, on whether how long it will take to find the global optimum, and on whether or not it will degenerate to simulated quenching. For the later use in the Simulated Annealing algorithm, let us define the new operator $\text{getTemperature}(t)$ which computes the temperature to be used at iteration t in the optimization process. For “getTemperature”, a few general rules hold. All schedules start with a temperature T_{start} which is greater than

zero. If the number of iterations t approaches infinity, the temperature must become 0K. This is a very weak statement, since we have shown that there exist finite boundaries after which Simulated Annealing is most likely to have converged. So there will be a finite t_{end} in all practical realizations after which the temperature drops to 0K, as shown in Equation 10.6.

$$T \in \mathbb{R}^+, t \in \mathbb{N}_0 \forall \quad T = \text{getTemperature}(t) \quad (10.3)$$

$$T_{start} = \text{getTemperature}(0) > 0 \quad (10.4)$$

$$\lim_{t \rightarrow \infty} \text{getTemperature}(t) = 0K \quad (10.5)$$

$$\exists t_{end} \in \mathbb{N} : \text{getTemperature}(t) = 0K \forall t \geq t_{end} \quad (10.6)$$

There exists a wide range of methods to determine this temperature schedule. Miki et al. [1331], for example, used genetic algorithms for this purpose. We will introduce only the three simple variants here given by Press et al. [1332].

- Reduce T to $(1 - \epsilon)T$ after every m iterations, where the exact values of $0 < \epsilon < 1$ and $m > 0$ are determined by experiment.
- Grant a total of K iterations, and reduce T after every m steps to a value $T = T_{start} \left(1 - \frac{t}{K}\right)^\alpha$ where t is the index of the current iteration and α is a constant, maybe 1, 2, or 4. α depends on the positions of the relative minima. Large values of α will spend more iterations at lower temperature.
- After every m moves, set T to β times $\Delta E_c = f(x_{cur}) - f(x^*)$, where β is an experimentally determined constant, $f(x_{cur})$ is the objective value of the currently examined solution candidate x_{cur} , and $f(x^*)$ is the objective value of the best phenotype x^* found so far. Since ΔE_c may be 0, we limit the temperature change to a maximum of $T * \gamma$ with $0 < \gamma < 1$.

If we let the temperature sink fast, we will lose the property of guaranteed convergence. In order to avoid getting stuck at local optima, we can then apply random restarting, which already has been discussed in the context of hill climbing in Section 8.5 on page 212.

10.4 Multi-Objective Simulated Annealing

Again, we want to combine this algorithm with multi-objective optimization and also enable it to return a set of optimal solutions. This can be done even simpler than in multi-objective hill climbing. Basically, we just need to replace the single objective function f with the fitness values v computed by a fitness assignment process on basis of the set of currently known best solutions (Arc), the currently investigated individual (p_{cur}), and the newly created points in the search space (p_{new}).

Algorithm 10.2: $x^* \leftarrow \text{simulatedAnnealingMO}(\text{cmp}_F, a)$

Input: f : the objective function to be minimized
Input: as : the maximum number of individuals allowed to be stored in the archive
Data: p_{new} : the newly generated individual
Data: p_{cur} : the point currently investigated in problem space
Data: Arc : the set of best individuals found so far
Data: T : the temperature of the system which is decreased over time
Data: t : the current time index
Data: ΔE : the energy difference of the x_{new} and x_{cur}
Output: x^* : the best element found

```

1 begin
2    $p_{new}.g \leftarrow \text{create}()$ 
   // Implicitly:  $p_{new}.x \leftarrow \text{gpm}(p_{new}.g)$ 
3    $p_{cur} \leftarrow p_{new}$ 
4    $Arc \leftarrow \text{createList}(1, p_{cur})$ 
5    $t \leftarrow 0$ 
6   while  $\text{terminationCriterion}()$  do
7      $v \leftarrow \text{assignFitness}(Arc \cup \{p_{new}, p_{cur}\}, \text{cmp}_F)$ 
8      $\Delta E \leftarrow v(p_{new}.x) - v(p_{cur}.x)$ 
9     if  $\Delta E \leq 0$  then
10       $p_{cur} \leftarrow p_{new}$ 
11    else
12       $T \leftarrow \text{getTemperature}(t)$ 
13      if  $\text{random}_u() < e^{-\frac{\Delta E}{k_B T}}$  then  $p_{cur} \leftarrow p_{new}$ 
14       $Arc \leftarrow \text{updateOptimalSet}(Arc, p_{new})$ 
15       $Arc \leftarrow \text{pruneOptimalSet}(Arc, as)$ 
16       $p_{new}.g \leftarrow \text{mutate}(p_{cur}.g)$ 
   // Implicitly:  $p_{new}.x \leftarrow \text{gpm}(p_{new}.g)$ 
17       $t \leftarrow t + 1$ 
18   return  $\text{extractPhenotypes}(Arc)$ 
19 end

```

Downhill Simplex (Nelder and Mead)

11.1 Introduction

The downhill simplex¹ (or Nelder-Mead method or amoeba algorithm²) published by Nelder and Mead [426] in 1965 is an single-objective optimization approach for searching the space of n -dimensional real vectors ($\mathbb{G} \subseteq \mathbb{R}^n$) [1333, 1334]. Historically, it is closely related to the simplex extension by Spendley et al. [425] to the Evolutionary Operation method mentioned in Section 2.1.2 on page 73 [427]. Since it only uses the values of the objective functions without any derivative information (explicit or implicit), it falls into the general class of direct search methods [1335, 1336], as most of the optimization approaches discussed in this book do.

Downhill simplex optimization uses $n+1$ points in the $\mathbb{R}^n$. These points form a polytope³, a generalization of a polygone, in the n -dimensional space – a line segment in $\mathbb{R}^1$, a triangle in $\mathbb{R}^2$, a tetrahedron in $\mathbb{R}^3$, and so on. Nondegenerated simplexes, i.e., those where the set of edges adjacent to any vertex form a basis in the $\mathbb{R}^n$, have one important feature: The result of replacing a vertex with its reflection through the opposite face is again, a nondegenerated simplex (see Fig. 11.1.1). The goal of downhill simplex optimization is to replace the best vertex of the simplex with an even better one or to ascertain that it is a candidate for the global optimum [427]. Therefore, its other points are constantly flipped around in an intelligent manner as we will outline in Section 11.3.

Like hill climbing approaches, the downhill simplex may not converge to the global minimum and can get stuck at local optima [1334, 1337, 1338]. Random restarts (as in Hill Climbing with Random Restarts discussed in Section 8.5 on page 212) can be helpful here.

11.2 General Information

11.2.1 Areas Of Application

Some example areas of application of downhill simplex are:

¹ http://en.wikipedia.org/wiki/Nelder-Mead_method [accessed 2008-06-14]

² In the book *Numerical Recipes in C++* by Press et al. [1332], this optimization method is called “amoeba algorithm”.

³ <http://en.wikipedia.org/wiki/Polytope> [accessed 2008-06-14]

Application	References
Chemistry, Chemical Engineering	[1339, 691, 689, 1340]
Robotics	[1341]
Physics	[1342, 1343]
Biochemistry	[1344]
Data Mining and Data Analysis	[1345]

11.2.2 Online Resources

Some general, online available resources on downhill simplex are:

<http://math.fullerton.edu/mathews/n2003/NelderMeadMod.html> [accessed 2008-06-14]

Last update: 2004-07-22

Description: Nelder-Mead Search for a Minimum

<http://www.boomer.org/c/p3/c11/c1106.html> [accessed 2008-06-14]

Last update: 2003

Description: Nelder-Mead (Simplex) Method

11.2.3 Books

Some books about (or including significant information about) downhill simplex are:

Avriel [1346]: *Nonlinear Programming: Analysis and Methods*

Walters, Morgan, Parker, Jr., and Deming [1339]: *Sequential Simplex Optimization: A Technique for Improving Quality and Productivity in Research, Development, and Manufacturing*

Press, Teukolsky, Vetterling, and Flannery [1332]: *Numerical Recipes in C++. Example Book. The Art of Scientific Computing*

11.3 The Downhill Simplex Algorithm

In Algorithm 11.1, we define the downhill simplex optimization approach. For simplification purposes we set both, the problem and the search space, to the n -dimensional real vectors, i. e., $\mathbb{X} \subseteq \mathbb{G} \subseteq \mathbb{R}^n$. In the actual implementation, we can use any set as problem space, given that a genotype-phenotype mapping $\text{gpm} : \mathbb{R}^n \mapsto \mathbb{X}$ is provided. Furthermore, notice that we optimize only a single objective function f . We can easily extend this algorithm for multi-objective optimization by using a comparator function cmp_F based on a set of objective functions F instead of comparing the values of f . In Algorithm 8.2, we have created a multi-objective hill climbing method with the same approach.

For visualization purposes, we apply the downhill simplex method exemplarily to an optimization problem with $\mathbb{G} = \mathbb{X} = \mathbb{R}^2$, where the simplex S consists of three points, in Figure 11.1.

The optimization process described by Algorithm 11.1 starts with creating a sample of $n + 1$ random points in the search space in line 2. Here, the `createPop` operation must ensure that these samples form a nondegenerated simplex. Notice that apart from the creation of the initial simplex, all further steps are deterministic and do not involve random numbers.

In each search step, the points in the simplex S are arranged in ascending order according to their corresponding objective values (line 4). Hence, the best solution candidate is $S_{[0]}$ and the worst is $S_{[n]}$. We then compute the center $\mathbf{m}$ of the n best points in line 5 and then

Algorithm 11.1: $x^* \leftarrow \text{downhillSimplex}(f)$

Input: f : the objective function subject to minimization
Input: [implicit] n : the dimension of the search space
Input: [implicit] $\alpha, \rho, \gamma, \sigma$: the reflection, the expansion, the contraction, and the shrink coefficient
Data: S : the simplex
Data: $\mathbf{m}$: the centroid of the simplex
Data: $\mathbf{r}$: the reflexion
Data: $\mathbf{e}$: the expansion
Data: $\mathbf{c}$: the contraction
Data: i : a counter variable
Output: x^* : the best solution candidate found

```

1 begin
2    $S \leftarrow \text{createPop}(n + 1)$ 
3   while  $\text{terminationCriterion}()$  do
4      $S \leftarrow \text{sortList}_a(S, f)$ 
5      $\mathbf{m} \leftarrow \frac{1}{n} \sum_{i=0}^{n-1} S[i]$ 
6     // Reflection: reflect the worst point over  $\mathbf{m}$ 
7      $\mathbf{r} \leftarrow \mathbf{m} + \alpha (\mathbf{m} - S[n])$ 
8     if  $f(S[0]) < f(\mathbf{r}) < f(S[n])$  then
9        $S[n] \leftarrow \mathbf{r}$ 
10    else
11      if  $f(\mathbf{r}) < f(S[0])$  then
12        // Expansion: try to search farther in this direction
13         $\mathbf{e} \leftarrow \mathbf{r} + \gamma (\mathbf{r} - \mathbf{m})$ 
14        if  $f(\mathbf{e}) < f(\mathbf{r})$  then  $S[n] \leftarrow \mathbf{e}$ 
15        else  $S[n] \leftarrow \mathbf{r}$ 
16      else
17         $b \leftarrow \text{true}$ 
18        if  $f(\mathbf{r}) \geq f(S[n-1])$  then
19          // Contraction: a test point between  $\mathbf{r}$  and  $\mathbf{m}$ 
20           $\mathbf{c} \leftarrow \rho \mathbf{r} + (1 - \rho) \mathbf{m}$ 
21          if  $f(\mathbf{c}) \leq f(\mathbf{r})$  then
22             $S[n] \leftarrow \mathbf{c}$ 
23             $b \leftarrow \text{false}$ 
24        if  $b$  then
25          // Shrink towards the best solution candidate  $S[0]$ 
26          for  $i \leftarrow n$  down to 1 do
27             $S[i] \leftarrow S[0] + \sigma (S[i] - S[0])$ 
28  return  $S[0]$ 
29 end
  
```

reflect the worst solution candidate $S[n]$ through this point in line 6, obtaining the new point $\mathbf{r}$ as also illustrated in Fig. 11.1.1. The reflection parameter α is usually set to 1.

In the case that $\mathbf{r}$ is somewhere in between of the points in the current simplex, i. e., neither better than $S[0]$ nor as worse as $S[n]$, we directly replace $S[n]$ with it. This simple move was already present in the first simplex algorithm defined by Spendley et al. [425]. The contribution of Nelder and Mead [426] was to turn the simplex search into an optimization algorithm by adding new options. These special operators were designed for speeding up the optimization process by deforming the simplex in way that they suggested would better adapt to the objective functions [427].




Fig. 11.1: One possible step of the downhill simplex algorithm applied to a problem in $\mathbb{R}^2$.

If $\mathbf{r}$ is better than the best solution candidate $S[0]$, one of these operators is to expand the simplex further into this promising direction (line 11). As sketched in Fig. 11.1.2, we obtain the point $\mathbf{e}$ with the expansion parameter γ set to 1. We now choose the better one of these two points in order to replace $S[n]$.

If $\mathbf{r}$ was no better than $S[n]$, we the simplex is contracted by creating a point $\mathbf{c}$ somewhere in between of $\mathbf{r}$ and $\mathbf{m}$ in line 17. In Fig. 11.1.3, the contraction parameter ρ was set to $\frac{1}{2}$. We substitute $S[n]$ with $\mathbf{c}$ only if $\mathbf{c}$ is better than $\mathbf{r}$.

When everything else fails, we shrink the whole simplex in line 23 by moving all points (except $S[0]$) into the direction of the current optimum $S[0]$. The shrinking parameter σ normally has the value $\frac{1}{2}$, as is the case in the example outlined in Fig. 11.1.4.

11.4 Hybridizing with the Downhill Simplex

Interestingly, there are some similarities between evolutionary algorithms and the downhill simplex. Takahama and Sakai [1347], for instance, argue that the downhill simplex can be considered as an evolutionary algorithm with special selection and reproduction operators. Each search step of Nelder and Mead's algorithm could be regarded as an n -ary reproduction operation for search spaces that are subsets the $\mathbb{R}^n$. Also, there are vague similarities between search operations of Differential Evolution (described in Section 5.5 on page 193), Particle Swarm Optimization (introduced in Chapter 14 on page 237), and the reflection operator of downhill simplex. The joint work of Jakumeit et al. [1348] and Barth et al. [1349], for example, goes into the direction of utilizing these similarities.

The research of Wang and Qiu [1350], [1351, 1352, 1353, 1354] focuses on the Hybrid Simplex Method PSO (HSMPSO), which, as the name says, is hybridization of Particle Swarm Optimization with Nelder and Mead's algorithm. In this approach, the downhill simplex operator is applied to each particle after a definite interval of iterations. Similar ideas of combining PSO with simplex methods are pursued by Fan et al. [1355], [1356].

Gao and Wang [1159] emphasize the close similarities between the reproduction operators of Differential Evolution and the search step of the downhill simplex. Thus, it seems only logical to combine the two methods in form of a new Memetic Algorithm. The Low Dimensional Simplex Evolution (LDSE) of Luo and Yu [1357] incorporates the single search steps of the downhill simplex applied to a number m of points which is lower than the actual dimensionality of the problem n . Luo and Yu [1357], [1358] reported that this method is able to outperform Differential Evolution when applied to the test function set of Ali et al. [1163].

There exist various other methods of hybridizing real-coded genetic algorithms with downhill simplex algorithm. Renders and Bersini [1359], for example, divide the population

into groups of $n + 1$ genotypes and allow the genetic algorithm to choose the Nelder-Mead simplex as additional reproduction operation besides crossover and simple averaging. The concurrent simplex of Yen et al. [1344; 1360], [1361] uses a probabilistic simplex method with $n + \Omega$ points instead of one, where the n best points are used to compute the centroid and the other $\Omega > 1$ points are reflected over it. They apply this idea to the top S individuals in the population, obtain $S - n$ children, and copy the best n individuals into the next generation. The remaining $(ps - S)$ genotypes (where ps is the population size) are created according to the conventional genetic algorithm reproduction scheme.

Barbosa et al. [1340] also use the downhill simplex as reproduction operator in a real-coded genetic algorithm. Their operator performs up to a given number of simplex search steps (20 in their work) and leads to improved results. Again, this idea goes more into the direction of Memetic Algorithms. Further approaches for hybridizing genetic algorithms with the downhill simplex have been contributed by Musil et al. [1343], Zhang et al. [1362], [1363], and Satapathy et al. [1345].

Tabu Search

*“I’d rather take the time to burn every last bridge I’ve ever crossed
beneath the sun
Than live my life knowin’ you may one day follow me over one.”*

— Aesop Rock [1364]

12.1 Introduction

Tabu Search¹ (TS) has been developed by Glover [1365] in the late 1980s. Basic ideas have also been introduced by Hansen [1366]. Further contributions in terms of formalizing this method have been made by Glover [1367; 1368], and de Werra and Hertz [1369] (as summarized by Hertz et al. [1370] in their tutorial on Tabu Search) as well as by Battiti and Tecchiolli [1371] and Cvijović and Klinowski [1372].

The word “tabu”² stems from Polynesia and describes a sacred place or object. Things that are *tabu* must be left alone and may not be visited or touched. Tabu Search extends hill climbing by this concept – it declares solution candidates which have already been visited as tabu. Hence, they must not be visited again and the optimization process is less likely to get stuck on a local optimum. The simplest realization of this approach is to use a list *tabu* which stores all solution candidates that have already been tested. If a newly created phenotype can be found in this list, it is not investigated but rejected right away. Of course, the list cannot grow infinitely but has a finite maximum length n . If the $n + 1^{st}$ solution candidate is added, the first one must be removed. Alternatively, this list could also be reduced with clustering. If some distance measure in the problem space $\mathbb{X}$ is available, a certain perimeter around the listed solution candidates can be declared as tabu. More complex approaches will store specific properties of the individuals instead of the phenotypes themselves in the list. This will not only lead to more complicated algorithms, but may also reject new solutions which are actually very good. Therefore, aspiration criteria will often be defined which override the tabu list and allow certain individuals.

¹ http://en.wikipedia.org/wiki/Tabu_search [accessed 2007-07-03]

² http://en.wikipedia.org/wiki/Tabu_%28Polynesian_culture%29 [accessed 2008-03-27]

Algorithm 12.1: $x^* \leftarrow \text{tabuSearch}(f, n)$

Input: f : the objective function subject to minimization
Input: n : the maximum length of the tabu list ($n > 0$)
Data: p_{new} : the new element created
Data: p^* : the (currently) best individual
Data: tabu : the tabu list
Output: x^* : the best element found

```

1 begin
2    $p^*.g \leftarrow \text{create}()$ 
   // Implicitly:  $p^*.x \leftarrow \text{gpm}(p^*.g)$ 
3    $\text{tabu} \leftarrow \text{createList}(1, p^*.x)$ 
4   while terminationCriterion() do
5      $p_{\text{new}}.g \leftarrow \text{mutate}(p^*.g)$ 
   // Implicitly:  $p_{\text{new}}.x \leftarrow \text{gpm}(p_{\text{new}}.g)$ 
6     if searchItemu( $p_{\text{new}}.x, \text{tabu}$ ) < 0 then
7       if  $f(p_{\text{new}}.x) < f(p^*.x)$  then  $p^* \leftarrow p_{\text{new}}$ 
8       if  $\text{len}(\text{tabu}) \geq n$  then  $\text{tabu} \leftarrow \text{deleteListItem}(\text{tabu}, 0)$ 
9        $\text{tabu} \leftarrow \text{addListItem}(\text{tabu}, p_{\text{new}}.x)$ 
10  return  $p^*.x$ 
11 end
  
```

12.2 General Information

12.2.1 Areas Of Application

Some example areas of application of Tabu Search are:

Application	References
Combinatorial Optimization	[1373, 1374, 1375, 1376, 1377, 1378, 1379, 405]
Machine Learning	[1380, 1369]
Biochemistry	[1381]
Operations Research	[1382]
Networking and Communication	[1383, 1384, 1385, 717, 1386]

12.2.2 Books

Some books about (or including significant information about) Tabu Search are:

Pardalos and Du [405]: <i>Handbook of Combinatorial Optimization</i>
Badiru [1330]: <i>Handbook of Industrial and Systems Engineering</i>
Reeves [406]: <i>Modern Heuristic Techniques for Combinatorial Problems</i>

12.3 Multi-Objective Tabu Search

The simple Tabu Search is very similar to hill climbing and Simulated Annealing, as you can see when comparing it with Chapter 8 on page 209 and Chapter 10 on page 217). With Algorithm 12.2, we thus can define a multi-objective variant for Tabu Search in a manner very similar to the multi-objective hill climbing or multi-objective Simulated Annealing.

Algorithm 12.2: $x^* \leftarrow \text{tabuSearchMO}(\text{cmp}_F, n, a)$

Input: cmp_F : the prevalence comparator
Input: n : the maximum length of the tabu list ($n > 0$)
Input: as : the maximum archive size
Data: tabu : the tabu list
Data: p_{new} : the new individual generated
Data: Arc : the set of best individuals known
Output: X^* : the set of the best elements found

```

1 begin
2    $\text{Arc} \leftarrow ()$ 
3    $p_{\text{new}}.x \leftarrow \text{create}()$ 
4   // Implicitly:  $p_{\text{new}}.x \leftarrow \text{gpm}(p_{\text{new}}.g)$ 
5    $\text{tabu} \leftarrow ()$ 
6   while  $\text{terminationCriterion}()$  do
7     if  $\text{searchItem}_u(p_{\text{new}}.x, \text{tabu} \cup \text{Arc}) < 0$  then
8        $\text{Arc} \leftarrow \text{updateOptimalSet}(\text{Arc}, p_{\text{new}})$ 
9        $\text{Arc} \leftarrow \text{pruneOptimalSet}(\text{Arc}, as)$ 
10       $v \leftarrow \text{assignFitness}(\text{Arc}, \text{cmp}_F)$ 
11      if  $\text{len}(\text{tabu}) \geq n$  then  $\text{tabu} \leftarrow \text{deleteListItem}(\text{tabu}, 0)$ 
12       $\text{tabu} \leftarrow \text{addListItem}(\text{tabu}, p_{\text{new}}.x)$ 
13       $p_{\text{new}} \leftarrow \text{select}(\text{Arc}, v, 1)[0]$ 
14       $p_{\text{new}}.g \leftarrow \text{mutate}(p_{\text{new}}.g)$ 
15      // Implicitly:  $p_{\text{new}}.x \leftarrow \text{gpm}(p_{\text{new}}.g)$ 
16  return  $\text{extractPhenotypes}(\text{Arc})$ 
17 end
  
```

Ant Colony Optimization

13.1 Introduction

Inspired by the research of Deneubourg et al. [1387], [1388, 1389] on real ants and probably by the simulation experiments of Stickland et al. [1390], Dorigo et al. [1391] developed the Ant Colony Optimization¹ (ACO) algorithm for problems that can be reduced to finding optimal paths in graphs in 1996. [1392, 1393, 1394, 1395] Ant Colony Optimization is based on the metaphor of ants that seek for food. An ant will leave the anthill and begin to wander randomly around. While the little insect paces around, it lays a trail of pheromone. Thus, after the ant has found some food, it can track its way back. By doing so, it distributes another layer of pheromone on the path. An ant that senses the pheromone will follow its trail with a certain probability. Each ant finding the food this way will excrete some pheromone on the path. By time, the pheromone density of the path will increase and more and more ants will follow it to the food and back. The higher the pheromone density, the more likely will an ant stay on a trail. However, the pheromones vaporize after some time. If all the food has been collected, they will no longer be renewed and the path will disappear after a while. Now, the ants will head to new, random locations.

This process of distributing and tracking pheromones is one form of stigmergy² and was first described by Grassé [1396]. Today, we subsume many different ways of communication by modifying the environment under this term, which can be divided into two groups: sematectonic and sign-based [1397]. According to Wilson [1398], we call modifications in the environment due to a task-related action which leads other entities involved in this task to change their behavior sematectonic stigmergy. If an ant drops a ball of mud somewhere, this may cause other ants to place mud balls at the same location. Step by step, these effects can cumulatively lead to the growth of complex structures. Sematectonic stigmergy has been simulated on computer systems by, for instance, Théraulaz and Bonabeau [1399] and with robotic systems by Werfel and Nagpal [1400], [1401, 1402].

The second form, sign-based stigmergy, is not directly task-related. It has been developed by social insects which use a wide range of pheromones and hormones for communication. Computer simulations for sign-based stigmergy were first performed by Stickland et al. [1390] in 1992.

The sign-based stigmergy is copied by Ant Colony Optimization [1391], where optimization problems are visualized as (directed) graphs. First, a set of *ants* performs randomized walks through the graphs. Proportional to the goodness of the solutions denoted by the paths, pheromones are laid out, i.e., the probability to walk into the direction of the paths is shifted. The ants run again through the graph, following the previously distributed pheromone. However, they will not exactly follow these paths. Instead, they may deviate from these

¹ http://en.wikipedia.org/wiki/Ant_colony_optimization [accessed 2007-07-03]

² <http://en.wikipedia.org/wiki/Stigmergy> [accessed 2007-07-03]

routes by taking other turns at junctions, for example, since their walk is still randomized. The pheromones just modify the probability distributions.

It is interesting to note that even real vector optimizations can be mapped to a graph problem, as introduced by Korošec and Šilc [1403]. Thanks to such ideas, the applicability of Ant Colony Optimization is greatly increased.

13.2 General Information

13.2.1 Areas Of Application

Some example areas of application of Ant Colony Optimization are:

Application	References
Combinatorial Optimization	[1404, 1405, 248, 1406, 1407]
Scheduling	[1408]
Networking and Communication	[1397, 1409, 1410, 1411, 1412, 1413, 1414, 1415, 1416, 1417] Section 23.1 on page 347

13.2.2 Conferences, Workshops, etc.

Some conferences, workshops and such and such on Ant Colony Optimization are:

<i>BIOMA</i> : International Conference on Bioinspired Optimization Methods and their Applications see Section 2.2.2 on page 77
<i>CEC</i> : Congress on Evolutionary Computation see Section 2.2.2 on page 77
<i>GECCO</i> : Genetic and Evolutionary Computation Conference see Section 2.2.2 on page 79
<i>ICNC</i> : International Conference on Advances in Natural Computation see Section 1.6.2 on page 62

13.2.3 Journals

Some journals that deal (at least partially) with Ant Colony Optimization are:

<i>Adaptive Behavior</i> , ISSN: Online: 1741-2633, Print: 1059-7123, appears quaterly, editor(s): Peter M. Todd, publisher: Sage Publications, http://www.isab.org/journal/ [accessed 2007-09-16], http://adb.sagepub.com/ [accessed 2007-09-16]
<i>Artificial Life</i> , ISSN: 1064-5462, appears quaterly, editor(s): Mark A. Bedau, publisher: MIT Press, http://www.mitpressjournals.org/loi/artl [accessed 2007-09-16]
<i>IEEE Transactions on Evolutionary Computation</i> (see Section 2.2.3 on page 80)
<i>The Journal of the Operational Research Society</i> (see Section 1.6.3 on page 64)

13.2.4 Online Resources

Some general, online available ressources on Ant Colony Optimization are:

<http://iridia.ulb.ac.be/~mdorigo/ACO/> [accessed 2007-09-13]

Last update: up-to-date

Description: Repository of books, publications, people, jobs, and software about ACO.

<http://uk.geocities.com/markcsinclair/aco.html> [accessed 2007-09-13]

Last update: 2006-11-17

Description: Small intro to ACO, some references, and a nice applet demonstrating its application to the travelling salesman problem.

Particle Swarm Optimization

14.1 Introduction

Particle Swarm Optimization¹ (PSO), developed by Eberhart and Kennedy [1418], [1419] in 1995, is a form of swarm intelligence in which the behavior of a biological social system like a flock of birds or a school of fish is simulated. When a swarm looks for food, its individuals will spread in the environment and move around independently. Each individual has a degree of freedom or randomness in its movements which enables it to find food accumulations. So, sooner or later, one of them will find something digestible and, being social, announces this to its neighbors. These can then approach the source of food, too. Each individual has a degree of freedom or randomness in its movements which enables it to find new food accumulations. Particle Swarm Optimization has been discussed, improved, and refined by many researchers such as Venter and Sobieszczanski-Sobieski [1420], Cai et al. [1421], Gao and Duan [1422], and Gao and Ren [1423]. Comparisons with other evolutionary approaches have been provided by Eberhart and Shi [1424] and Angeline [1425].

With Particle Swarm Optimization, a swarm of particles (individuals) in a n -dimensional search space $\mathbb{G}$ is simulated, where each particle p has a position $p.g \in \mathbb{G} \subseteq \mathbb{R}^n$ and a velocity $p.v \in \mathbb{R}^n$. The position $p.g$ corresponds to the genotypes, and, in most cases, also to the solution candidates, i. e., $p.x = p.g$, since most often the problem space $\mathbb{X}$ is also the $\mathbb{R}^n$ and $\mathbb{X} = \mathbb{G}$. However, this is not necessarily the case and generally, we can introduce any form of genotype-phenotype mapping in Particle Swarm Optimization. The velocity vector $p.v$ of an individual p determines in which direction the search will continue and if it has an explorative (high velocity) or an exploitive (low velocity) character.

In the initialization phase of Particle Swarm Optimization, the positions and velocities of all individuals are randomly initialized. In each step, first the velocity of a particle is updated and then its position. Therefore, each particle p has a memory holding its best position $\text{best}(p) \in \mathbb{G}$. In order to realize the social component, the particle furthermore knows a set of topological neighbors $N(p)$. This set could be defined to contain adjacent particles within a specific perimeter, i. e., all individuals which are no further away from $p.g$ than a given distance δ according to a certain distance measure² dist . Using the Euclidian distance measure $\text{dist}_{\text{eucl}}$ specified in Definition 28.8 on page 470 we get:

$$\forall p, q \in \text{Pop} : q \in N(p) \Leftrightarrow \text{dist}_{\text{eucl}}(p.g, q.g) \leq \delta \quad (14.1)$$

Each particle can communicate with its neighbors, so the best position found so far by any element in $N(p)$ is known to all of them as $\text{best}(N(p))$. The best position ever visited by any individual in the population (which the optimization algorithm always keeps track of) is $\text{best}(\text{Pop})$.

¹ http://en.wikipedia.org/wiki/Particle_swarm_optimization [accessed 2007-07-03]

² See Section 28.1 on page 469 for more information on distance measures.

The PSO algorithm may make use of either $\text{best}(N(p))$ or $\text{best}(Pop)$ for adjusting the velocity of the particle p . If it relies on the global best position, the algorithm will converge fast but may find the global optimum less probably. If, on the other hand, neighborhood communication is used, the convergence speed drops but the global optimum is found more likely.

Definition 14.1 (psoUpdate). The search operation $q = \text{psoUpdate}(p, Pop)$ applied in Particle Swarm Optimization creates a new particles q to replace an existing one (p) by incorporating its genotype $p.g$, its velocity $p.v$. We distinguish local updating (Equation 14.3) and global updating (Equation 14.2), which additionally uses the data from the whole population Pop . psoUpdate thus fulfills one of these two equations and Equation 14.4, showing how the i^{th} components of the corresponding vectors are computed.

$$q.v_i = p.v_i + (\text{random}_u(0, \mathbf{c}_i) * (\text{best}(p).g_i - p.g_i)) + (\text{random}_u(0, \mathbf{d}_i) * (\text{best}(Pop).g_i - p.g_i)) \quad (14.2)$$

$$q.v_i = p.v_i + (\text{random}_u(0, \mathbf{c}_i) * (\text{best}(p).g_i - p.g_i)) + (\text{random}_u(0, \mathbf{d}_i) * (\text{best}(N(p)).g_i - p.g_i)) \quad (14.3)$$

$$q.g_i = p.g_i + p.v_i \quad (14.4)$$

The learning rate vectors $\mathbf{c}$ and $\mathbf{d}$ have strong influence of the convergence speed of Particle Swarm Optimization. The search space $\mathbb{G}$ (and thus, also the values of $p.g$) is normally confined by minimum and maximum boundaries. For the absolute values of the velocity, normally maximum thresholds also exist. Thus, real implementations of “psoUpdate” have to check and refine their results before the utility of the solution candidates is evaluated.

Algorithm 14.1 illustrates the native form of the Particle Swarm Optimization using the update procedure from Definition 14.1. Like hill climbing, this algorithm can easily be generalized for multi-objective optimization and for returning sets of optimal solutions (compare with Section 8.3 on page 210).

Algorithm 14.1: $x^* \leftarrow \text{psoOptimizer } f \text{ } ps$

Input: f : the function to optimize

Input: ps : the population size

Data: Pop : the particle population

Data: i : a counter variable

Output: x^* : the best value found

```

1 begin
2    $Pop \leftarrow \text{createPop}(ps)$ 
3   while  $\text{terminationCriterion}()$  do
4     for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
5        $Pop[i] \leftarrow \text{psoUpdate}(Pop[i], Pop)$ 
6   return  $\text{best}(Pop).x$ 
7 end
```

14.2 General Information

14.2.1 Areas Of Application

Some example areas of application of particle swarm optimization are:

Application	References
Machine Learning	[1419, 1426, 1427]
Function Optimization	[1419, 1428]
Geometry and Physics	[1429]
Operations Research	[1430]
Chemistry, Chemical Engineering	[1431, 1432]

14.2.2 Online Resources

Some general, online available resources on particle swarm optimization are:

http://www.swarmintelligence.org/ [accessed 2007-08-26]
Last update: up-to-date
Description: Particle Swarm Optimization Website by Xiaohui Hu
http://www.red3d.com/cwr/boids/ [accessed 2007-08-26]
Last update: up-to-date
Description: Boids – Background and Update by Craig Reynolds
http://www.projectcomputing.com/resources/psovis/ [accessed 2007-08-26]
Last update: 2004
Description: Particle Swarm Optimization (PSO) Visualisation (or “PSO Visualization”)
http://www.engr.iupui.edu/~eberhart/ [accessed 2007-08-26]
Last update: 2003
Description: Russ Eberhart’s Home Page
http://www.cis.syr.edu/~mohan/psa/ [accessed 2007-08-26]
Last update: 1999
Description: Particle Swarm Optimization Homepage
http://tracer.uc3m.es/tws/psa/ [accessed 2007-11-06]
Last update: up-to-date
Description: Website on Particle Swarm Optimization

14.2.3 Conferences, Workshops, etc.

Some conferences, workshops and such and such on particle swarm optimization are:

<i>GECCO</i> : Genetic and Evolutionary Computation Conference
see Section 2.2.2 on page 79
<i>ICNC</i> : International Conference on Advances in Natural Computation
see Section 1.6.2 on page 62
<i>SIS</i> : IEEE Swarm Intelligence Symposium
http://www.computelligence.org/sis/ [accessed 2007-08-26]
History: 2007: Honolulu, Hawaii, USA, see [1433]
2006: Indianapolis, IN, USA, see [1434]
2005: Pasadena, CA, USA, see [1435]
2003: Indianapolis, IN, USA, see [1436]

14.2.4 Books

Some books about (or including significant information about) particle swarm optimization are:

Memetic Algorithms

Memetic Algorithms¹ (MAs) are a family of optimization methods introduced by Moscato [1438]. Memetic Algorithms try to simulate cultural evolution rather than the biological one. In order to point out the differences to evolutionary algorithms, Moscato [1438] uses the example of Chinese martial art Kung-Fu which has developed over many generations of masters teaching their students certain sequences of movements, the so-called *forms*. Each form is composed of a set of elementary aggressive and defensive patterns. Dawkins [1439] has defined the concept *meme* as “a unit of imitation in cultural transmission” which here fits exactly. In martial arts, each of the undecomposable sub-movements of a form represents one meme. New memes are only seldom introduced and only few amongst the masters of the art have the ability to do so. Also, only the best of the population of Kung-Fu practitioners can become masters and teach disciples.

Memetic Algorithms are a combination of population-based global optimization and heuristic local search [1438, 1440, 1441, 1442, 1443, 1444]. Similar to evolutionary algorithms, the population is initialized with random individuals. Starting with each single individual, a local search is performed corresponding to the improvements in skill a Kung-Fu fighter makes during his training. After that, the individuals start to interact

- competitively in the form of selection corresponding to Kung-Fu tournaments or
- cooperatively in the form of recombination corresponding to the exchange of knowledge between two Kung-Fu masters.

This alternation between local refinement and global selection and reproduction is repeated until a stopping criterion is met.

First experiences in hybridizing genetic algorithms with other search algorithms have been reported by Goldberg [139]. In contrast with genetic algorithms, the elements of the genotypes g are considered to be memes instead of genes. Comparisons of Memetic Algorithms with other optimization problems have been provided by Merz and Freisleben [1445]. The definition of Memetic Algorithms given by Moscato [1438] is relatively general and encompass many different approaches.

In the context of this work, we will regard them as two-level optimization algorithms, where the top-level algorithm is an evolutionary algorithm or otherwise population-based algorithm and at the bottom-level, a single individual optimizer like hill climbing, Simulated Annealing, the downhill simplex (see Section 11.4) or an informed search are applied.

In order to create such algorithms, one would replace the calls to `createPop` and `reproducePop` in the basic evolutionary algorithm scheme `simpleEA` defined in Algorithm 2.1 on page 72 with the refined operations `createPopMA` and `reproducePopMA` from Memetic Algorithms. These two extensions, introduced in Algorithm 15.1 and Algorithm 15.2 use a local search operator “`optimizeLocal`” which stands for the application of any of the previously mentioned single-individual optimizers.

¹ http://en.wikipedia.org/wiki/Memetic_algorithm [accessed 2007-07-03]

Algorithm 15.1: $Pop \leftarrow \text{createPopMA}(s, \text{cmp}_F)$

Input: s : the size of the population to be created
Input: cmp_F : the optimization criterion
Data: i : a counter variable
Data: p : an individual record
Output: Pop : the new, randomly created but locally refined population

```

1 begin
2    $Pop \leftarrow ()$ 
3   for  $i \leftarrow 1$  up to  $s$  do
4      $p \leftarrow \text{create}()$ 
5      $p \leftarrow \text{optimizeLocal}(p, \text{cmp}_F)$ 
6      $Pop \leftarrow \text{addListItem}(Pop, p)$ 
7   return  $Pop$ 
8 end

```

Algorithm 15.2: $Pop \leftarrow \text{reproducePopMA}(Mate, \text{cmp}_F)$

Input: $Mate$: the mating pool
Input: cmp_F : the optimization criterion
Data: i : a counter variable
Data: p : an individual record
Output: Pop : the population containing the locally optimized offspring of $Mate$

```

1 begin
2    $Pop \leftarrow \text{reproducePop}(Mate)$ 
3   for  $i \leftarrow 0$  up to  $\text{len}(Pop) - 1$  do
4      $p \leftarrow Pop[i]$ 
5      $p \leftarrow \text{optimizeLocal}(p, \text{cmp}_F)$ 
6      $Pop[i] \leftarrow p$ 
7   return  $Pop$ 
8 end

```

15.1 General Information

15.1.1 Areas Of Application

Some example areas of application of Memetic Algorithms are:

Application	References
Combinatorial Optimization	[1442, 1443, 1445, 1446]
Engineering, Structural Optimization, and Design	[1446]
Biochemistry	[1446]
Networking and Communication	[1447, 1448]
Scheduling	[1446]

15.1.2 Online Resources

Some general, online available ressources on Memetic Algorithms are:

http://www.densis.fee.unicamp.br/~moscato/memetic_home.html [accessed 2008-04-03]

Last update: 2002-08-16

Description: The Memetic Algorithms' Home Page by Pablo Moscato

15.1.3 Books

Some books about (or including significant information about) Memetic Algorithms are:

Hart, Krasnogor, and Smith [1446]: *Recent Advances in Memetic Algorithms*

State Space Search

16.1 Introduction

State space search strategies¹ are not directly counted to the optimization algorithms. In global optimization, objective functions $f \in F$ are optimized. Normally, all elements of the problem space $\mathbb{X}$ are valid and only differ in their *utility* as solution. Optimization is about to find the elements with the best utility. State space search is instead based on a criterion $\text{isGoal} : \mathbb{X} \mapsto \mathbb{B}$ which determines whether an element of the problem space is a valid solution or not. The purpose of the search process is to find elements x from the solution space $\mathbb{S} \subseteq \mathbb{X}$, i. e., those for which $\text{isGoal}(x)$ is **true**. [1276, 1449, 1450]

Definition 16.1 (*isGoal*). The function $\text{isGoal} : \mathbb{X} \mapsto \mathbb{B}$ is the target predicate of state space search algorithms which states whether a given state $x \in \mathbb{X}$ is a valid solution (by returning **true**), i. e., the goal state, or not (by returning **false**). isGoal thus corresponds to membership in the solution space $\mathbb{S}$ defined in Definition 1.20 on page 23.

$$x \in \mathbb{S} \Leftrightarrow \text{isGoal}(x) = \text{true} \quad (16.1)$$

We will be able to apply state space search strategies for global optimization if we can define a threshold $\tilde{y}_i$ for each objective function $f_i \in F$. If we assume minimization, $\text{isGoal}(x)$ becomes **true** if and only if the values of all objective functions for x drop below given thresholds, and thus, we can define isGoal according to Equation 16.2.

$$\text{isGoal}(x) = \begin{cases} \text{true} & \text{if } f_i(x) \leq \tilde{y}_i \ \forall i \in 1..|F| \\ \text{false} & \text{otherwise} \end{cases} \quad (16.2)$$

In state space search algorithms, the search space $\mathbb{G}$ and the problem space $\mathbb{X}$ are often identical. Most state space search can only be applied if the search space is enumerable. One feature of the search algorithms introduced here is that they all are deterministic. This means that they will yield the same results in each run (when applied to the same problem, that is).

One additional operator needs to be defined for state space search algorithms: the “expand” function which helps enumerating the search space.

Definition 16.2 (*expand*). The operator $\text{expand} : \mathbb{G} \mapsto \mathcal{P}(\mathbb{G})$ receives one element g from the search space $\mathbb{G}$ as input and computes a set G of elements which can be reached from it.

“expand” is the exploration operation of state space search algorithms. Different from the mutation operator of evolutionary algorithms, it is strictly deterministic and returns a set instead of single individual. Applying it to the same element g values will thus always

¹ http://en.wikipedia.org/wiki/State_space_search [accessed 2007-08-06]

yield the same set G . We can consider $\text{expand}(g)$ to return the set of all possible results that we could obtain with a unary search operation like mutation.

The realization of expand has severe impact on the performance of search algorithms. An efficient implementation, for example, should not include states that have already been visited in the returned set. If the same elements are returned, the same solution candidates and all of their children will be evaluated multiple times, which would be useless and time consuming. Another problem occurs if there are two elements $g_1, g_2 \in \mathbb{G}$ with $g_1 \in \text{expand}(g_2)$ and $g_2 \in \text{expand}(g_1)$ exist. Then, the search would get trapped in an endless loop. Thus, visiting a genotype twice should always be avoided. Often, it is possible to design the search operations in a way preventing this from the start. Otherwise, tabu lists should be used, as done in the previously discussed Tabu Search algorithm (see Chapter 12 on page 229).

Since we want to keep our algorithm definitions as general as possible, we will keep the notation of individuals p that encompass a genotype $p.g \in \mathbb{G}$ and a phenotype $p.x \in \mathbb{X}$. Therefore, we need to an expansion operator that returns a set of individuals P rather than a set G of elements of the search space. We therefore define the operation “ expandToInd ” in Algorithm 16.1.

Algorithm 16.1: $P \leftarrow \text{expandToInd}(g)$

Input: g : the element of the search space to be expanded
Data: i : a counter variable
Data: p : an individual record
Output: P : the list of individual records resulting from the expansion

```

1 begin
2    $G \leftarrow \text{expand}(g)$ 
3    $P \leftarrow ()$ 
4   for  $i \leftarrow 0$  up to  $\text{len}(G) - 1$  do
5      $p.g \leftarrow G[i]$ 
6     // Implicitly:  $p^*.x \leftarrow \text{gpm}(p^*.g)$ 
7      $P \leftarrow \text{addListItem}(P, p)$ 
8   return  $P$ 
9 end

```

For all state space search strategies, we can define four criteria that tell if they are suitable for a given problem or not.

1. *Completeness.* Does the search algorithm guarantee to find a solution (given that there exists one)? (Do not mix up with the completeness of search operations specified in Definition 1.27 on page 25.)
2. *Time Consumption.* How much time will the strategy need to find a solution?
3. *Memory Consumption.* How much memory will the algorithm need to store intermediate steps? Together with time consumption this property is closely related to complexity theory, as discussed in Section 29.1.3 on page 482.
4. *Optimality.* Will the algorithm find an optimal solution if there exist multiple correct solutions?

Search algorithms can further be classified according to the following definitions:

Definition 16.3 (Local Search). Local search algorithms work on a single current state (instead of multiple solution candidates) and generally transcend only to neighbors of the current state [1276].

Local search algorithms are not systematic but have two major advantages: They use very little memory (normally only a constant amount) and are often able to find solutions

in large or infinite search spaces. These advantages come, of course, with large trade-offs in processing time.

We can consider local searches as special case of global searches which incorporate larger populations. If the previously mentioned requirements are met, global search, in turn, can be regarded as a special case of global optimization algorithms.

16.2 General Information

16.2.1 Areas Of Application

Some example areas of application of State space search are:

Application	References
Networking and Communication	[1451, 1452, 1453]

16.2.2 Books

Some books about (or including significant information about) State space search are:

Russell and Norvig [1276]: <i>Artificial Intelligence: A Modern Approach</i>

16.3 Uninformed Search

The optimization algorithms that we have considered up to now always require some sort of utility measure. These measures, the objective functions, are normally real-valued and allow us to make fine distinctions between different individuals. Under some circumstances, however, only a criterion isGoal is given as a form of Boolean objective function. The methods previously discussed will then not be able to descend a gradient anymore and degenerate to random walks (see Section 16.3.5 on page 250).

Here, uninformed search strategies² are a viable alternative since they do not require or take into consideration any knowledge about the special nature of the problem (apart from the knowledge represented by the expand operation, of course). Such algorithms are very general and can be applied to a wide variety of problems. Their common drawback is that search spaces are often very large. Without the incorporation of information in form of heuristic functions, for example, the search may take very long and quickly becomes infeasible [1276, 1449, 1450].

16.3.1 Breadth-First Search

In breadth-first search³ (BFS), we start with expanding the root solution candidate. Then all of the states derived from this expansion are visited, and then all their children, and so on. In general, we first expand all states in depth d before considering any state in depth $d + 1$.

It is complete, since it will always find a solution if there exists one. If so, it will also find the solution that can be reached from the root state with the least expansion steps. Hence, if the number of expansion steps needed from the origin to a state is a measure for the costs, BFS is also optimal.

² http://en.wikipedia.org/wiki/Uninformed_search [accessed 2007-08-07]

³ http://en.wikipedia.org/wiki/Breadth-first_search [accessed 2007-08-06]

Algorithm 16.2: $X^* \leftarrow \text{bfs}(r, \text{isGoal})$

Input: r : the root individual to start the expansion at
Input: isGoal : an operator that checks whether a state is a goal state or not
Data: p : the state currently processed
Data: P : the queue of states to explore
Output: X^* : the solution states found, or $\emptyset$

```

1 begin
2    $P \leftarrow \text{createList}(1, r)$ 
3   while  $\text{len}(P) > 0$  do
4      $p \leftarrow \text{deleteListItem}(P, 0)$ 
5     if  $\text{isGoal}(p.x)$  then return  $\{p.x\}$ 
6      $P \leftarrow \text{appendList}(P, \text{expandToInd}(p.g))$ 
7   return  $\emptyset$ 
8 end

```

Algorithm 16.2 illustrates how breadth-first search works. The algorithm is initialized with a root state $r \in \mathbb{G}$ which marks the starting point of the search. BFS uses a state list P which initially only contains this root individual. In a loop, the first element p is removed from this list. If the goal predicate $\text{isGoal}(p.x)$ evaluates to **true**, $p.x$ is a goal state and we can return a set $X^* = \{p.x\}$ containing it as the solution. Otherwise, we expand $p.g$ and append the newly found individuals to the end of queue P . If no solution can be found, this process will continue until the whole accessible search space has been enumerated and P becomes empty. Then, an empty set is returned in place of X^* , because there is no element x in the (accessible part of the) problem space $\mathbb{X}$ for which $\text{isGoal}(x)$ becomes **true**.

In order to examine the space and time complexity of BFS, we assume a hypothetical state space $\mathbb{G}_h$ where the expansion of each state $g \in \mathbb{G}_h$ will return a set of $\text{len}(\text{expand}(g)) = b$ new states. In depth 0, we only have one state, the root state r . In depth 1, there are b states, and in depth 2 we can expand each of them to again, b new states which makes b^2 , and so on. Up to depth d we have a number of states total of

$$1 + b + b^2 + \dots + b^d = \frac{b^{d+1} - 1}{b - 1} \in \mathbf{O}(b^d) \quad (16.3)$$

We have both, a space and time complexity from $\mathbf{O}(b^d)$. In the worst case, all nodes in depth d need to be stored, in the best case only those of depth $d - 1$.

16.3.2 Depth-First Search

Depth-first search⁴ (DFS) is very similar to BFS. From the algorithmic point of view, the only difference that it uses a stack instead of a queue as internal storage for states (compare line 4 in Algorithm 16.3 with line 4 in Algorithm 16.2). Here, always the last state element of the set of expanded states is considered next. Thus, instead of searching level for level in the breath as BFS does, DFS searches in depth (which – believe it or not – is the reason for its name). DFS advances in depth until the current state cannot further be expanded, i. e., $\text{expand}(p.g) = ()$. Then, the search steps again up one level. If the whole search space has been browsed and no solution is found, $\emptyset$ is returned.

The memory consumption of the DFS is linear, because in depth d , at most $d * b$ states are held in memory. If we assume a maximum depth m , the time complexity is b^m in the worst case where the solution is the last child state in the path explored the last. If m is very large or infinite, a DFS may take very long to discover a solution or will not find it at all, since it may get stuck in a “wrong” branch of the state space. Hence, depth first search is neither complete nor optimal.

⁴ http://en.wikipedia.org/wiki/Depth-first_search [accessed 2007-08-06]

Algorithm 16.3: $X^* \leftarrow \text{dfs}(r, \text{isGoal})$

Input: r : the root individual to start the expansion at
Input: isGoal : an operator that checks whether a state is a goal state or not
Data: p : the state currently processed
Data: P : the queue of states to explore
Output: X^* : the solution states found, or $\emptyset$

```

1 begin
2    $P \leftarrow \text{createList}(1, r)$ 
3   while  $\text{len}(P) > 0$  do
4      $p \leftarrow \text{deleteListItem}(P, \text{len}(P) - 1)$ 
5     if  $\text{isGoal}(p.x)$  then return  $\{p.x\}$ 
6      $P \leftarrow \text{appendList}(P, \text{expandToInd}(p.g))$ 
7   return  $\emptyset$ 
8 end

```

16.3.3 Depth-limited Search

The depth-limited search⁵ [1276] is a depth-first search that only proceeds up to a given maximum depth d . In other words, it does not examine solution candidates that are more than d expand-operations away from the root state r , as outlined in Algorithm 16.4 in a recursive form. Analogously to the plain depth first search, the time complexity now becomes b^d and the memory complexity is in $\mathbf{O}(b * d)$. Of course, the depth-limited search can neither be complete nor optimal. If a maximum depth of the possible solutions however known, it may be sufficient.

Algorithm 16.4: $X^* \leftarrow \text{dl\_dfs}(r, \text{isGoal}, d)$

Input: r : the root individual to start the expansion at
Input: isGoal : an operator that checks whether a state is a goal state or not
Input: d : the (remaining) allowed depth steps
Data: p : the state currently processed
Output: X^* : the solution states found, or $\emptyset$

```

1 begin
2   if  $\text{isGoal}(r.x)$  then return  $\{r.x\}$ 
3   if  $d > 0$  then
4     foreach  $p \in \text{expandToInd } r.g$  do
5        $X^* \leftarrow \text{dl\_dfs}(p, \text{isGoal}, d - 1)$ 
6       if  $\text{len}(X^*) > 0$  then return  $X^*$ 
7   return  $\emptyset$ 
8 end

```

16.3.4 Iterative Deepening Depth-First Search

The iterative deepening depth-first search⁶ (IDDFS, [1276]), defined in Algorithm 16.5, iteratively runs a depth-limited DFS with stepwise increasing maximum depths d . In each iteration, it visits the states in the according to the depth-first search. Since the maximum

⁵ http://en.wikipedia.org/wiki/Depth-limited_search [accessed 2007-08-07]

⁶ <http://en.wikipedia.org/wiki/IDDFS> [accessed 2007-08-08]

depth is always incremented by one, one new level in terms means of distance in expand-operations from the root is explored in each iteration. This effectively leads to some form of breadth-first search.

IDDFS thus unites the advantages of BFS and DFS: It is complete and optimal, but only has a linearly rising memory consumption in $\mathbf{O}(d * b)$. The time consumption, of course, is still in $\mathbf{O}(b^d)$. IDDFS is the best uninformed search strategy and can be applied to large search spaces with unknown depth of the solution.

The Algorithm 16.5 is intended for infinitely large search spaces. In real systems, there is a maximum $\hat{d}$ after which the whole space would be explored and the algorithm should return $\emptyset$ if no solution was found.

Algorithm 16.5: $X^* \leftarrow \text{iddfs}(r, \text{isGoal})$

Input: r : the root individual to start the expansion at

Input: isGoal : an operator that checks whether a state is a goal state or not

Data: d : the current depth limit

Output: X^* : the solution states found, or $\emptyset$

```

1 begin
2    $d \leftarrow 0$ 
3   repeat
4      $X^* \leftarrow \text{dl\_dfs}(r, \text{isGoal}, d)$ 
5      $d \leftarrow d + 1$ 
6   until  $\text{len}(X^*) > 0$ 
7   return  $X^*$ 
8 end
```

16.3.5 Random Walks

Random walks⁷ (sometimes also called drunkard's walk) are a special case of undirected, local search. Instead of proceeding according to some schema like depth-first or breadth-first, the next solution candidate to be explored is always generated randomly from the currently investigated one. [1454, 1455] Under some special circumstances, random walks can be the search algorithms of choice. This for instance the case in

- If we encounter a state explosion because there are too many states to which we can possibly transcend to and methods like breadth-first search or iterative deepening depth-first search cannot be applied because they would consume too much memory.
- In certain cases of online search it is not possible to apply systematic approaches like BFS or DFS. If the environment, for instance, is only partially observable and each state transition represents an immediate interaction with this environment, we are maybe not able to navigate to past states again. One example for such a case is discussed in the work of Skubch [42; 43] about reasoning agents.

Random walks are often used in optimization theory for determining features of a fitness landscape. Measures that can be derived mathematically from a walk model include estimates for the number of steps needed until a certain configuration is found, the chance to find a certain configuration at all, and the average difference of the objective values of two consecutive populations. From practically running random walks, some information about the search space can be extracted. Skubch [42; 43], for instance, uses the number of encounters of certain state transition during the walk in order to successively build a heuristic function.

⁷ http://en.wikipedia.org/wiki/Random_walk [accessed 2007-11-27]

16.4 Informed Search

In an informed search⁸, a heuristic function h helps to decide which states are to be expanded next. If the heuristic is good, informed search algorithms may dramatically outperform uninformed strategies [9, 10, 11].

As specified in Definition 1.2 on page 4, heuristic functions are problem domain dependent. In the context of an informed search, a heuristic function $h : \mathbb{X} \mapsto \mathbb{R}^+$ maps the states in the state space $\mathbb{G}$ to the positive real numbers $\mathbb{R}^+$. We further define that all heuristics will be zero for the elements which are part of the solution space $\mathbb{S}$, i. e.,

$$\forall x \in \mathbb{X} : \text{isGoal}(x) \Rightarrow h(x) = 0 \quad \forall \text{heuristics } h : \mathbb{X} \mapsto \mathbb{R}^+ \quad (16.4)$$

There are two possible meanings of the values returned by a heuristic function h :

1. In the above sense, the value of a heuristic function $h(p.x)$ for an individual p is the higher, the more expand-steps $p.g$ is *probably* (or *approximately*) away from finding a valid solution. Hence, the heuristic function represents the distance of an individual to a solution in *solution space*.
2. The heuristic function can also represent an objective function in some way. Suppose that we know the minimal value $\tilde{y}$ for an objective function f or at least a value from where on all solutions are feasible. If this is the case, we could set $h(p.x) = \max\{0, f(p.x) - \tilde{y}\}$, assuming that f is subject to minimization. Now the value of heuristic function will be the smaller, the closer an individual is to a possible correct solution and Equation 16.4 still holds. In other words, a heuristic function may also represent the distance to a solution in *objective space*.

Of course, both meanings are often closely related since states that are close to each other in problem space are probably also close to each other in objective space (the opposite does not necessarily hold).

A best-first search⁹ [11] is a search algorithm that incorporates such a heuristic function h in a way which ensures that promising individuals p with low estimation values $h(p.x)$ are evaluated before other states q that receive a higher values $h(q.x) > h(p.x)$.

16.4.1 Greedy Search

“And what’s mines is mines and yours is mine”

— Sticky Fingaz [1456]

A greedy search¹⁰ is a best-first search where the currently known solution candidate with the lowest heuristic value is investigated next. The greedy algorithm internal sorts the list of currently known states in *descending* order according to a heuristic function h . Thus, the elements with the best (lowest) heuristic value will be at the end of the list, which then can be used as a stack. The greedy search as specified in Algorithm 16.6 now works like a depth-first search on this stack and thus, also shares most of the properties of the DFS. It is neither complete nor optimal and its worst case time consumption is b^m . On the other hand, like breadth-first search, its worst-case memory consumption is also b^m .

⁸ http://en.wikipedia.org/wiki/Search_algorithms#Informed_search [accessed 2007-08-08]

⁹ http://en.wikipedia.org/wiki/Best-first_search [accessed 2007-09-25]

¹⁰ http://en.wikipedia.org/wiki/Greedy_search [accessed 2007-08-08]

Algorithm 16.6: $X^* \leftarrow \text{greedySearch}(r, \text{isGoal}, h)$

Input: r : the root individual to start the expansion at
Input: isGoal : an operator that checks whether a state is a goal state or not
Input: h : the heuristic function
Data: p : the state currently processed
Data: P : the queue of states to explore
Output: X^* : the solution states found, or $\emptyset$

```

1 begin
2    $P \leftarrow \text{createList}(1, r)$ 
3   while  $\text{len}(P) > 0$  do
4      $P \leftarrow \text{sortList}_d(P, \text{cmp}(p_1, p_2) \equiv h(p_1.x) - h(p_2.x))$ 
5      $p \leftarrow \text{deleteListItem}(P, \text{len}(P) - 1)$ 
6     if  $\text{isGoal}(p.x)$  then return  $\{p.x\}$ 
7      $P \leftarrow \text{appendList}(P, \text{expandToInd}(p.g))$ 
8   return  $\emptyset$ 
9 end

```

16.4.2 A* search

In A* search¹¹ is a best-first search that uses a estimation function $h^* : \mathbb{X} \mapsto \mathbb{R}^+$ which is the sum of a heuristic function $h(x)$ that estimates the costs needed to get from x to a valid solution and a function $g(x)$ that computes the costs that are needed to get to x .

$$h^*(x) = g(x) + h(x) \quad (16.5)$$

A* search proceeds exactly like the greedy search outlined in Algorithm 16.6, if h^* is used instead of plain h . An A* search will definitely find a solution if there exists one, i. e., it is complete.

Definition 16.4 (Admissible Heuristic Function). A heuristic function $h : \mathbb{X} \mapsto \mathbb{R}^+$ is admissible if it never overestimates the minimal costs for reaching a goal state.

Definition 16.5 (Monotonic Heuristic Function). A heuristic function $h : \mathbb{X} \mapsto \mathbb{R}^+$ is monotonic¹² if it never overestimates the costs for getting from one state to its successor.

$$h(p.x) \leq g(q.x) - g(p.x) + h(q.x) \quad \forall q.g \in \text{expand}(p.g) \quad (16.6)$$

An A* search is optimal if the heuristic function h used is admissible. Optimal in this case means that there exists no search algorithm that can find the same solution as the A* search needing fewer expansion steps if using the same heuristic. If we implement expand in a way which prevents that a state is visited more than once, h also needs to be monotone in order for the search to be optimal.

16.4.3 Adaptive Walks

An *adaptive walk* is a theoretical optimization method which, like a random walk, usually works on a population of size 1. It starts at a random location in the search space and proceeds by changing (or mutating) its single solution candidate. For this modification, three methods are available:

- *One-mutant change:* The optimization process chooses a single new individual from the set of “one-mutant change” neighbors, i. e., a neighboring individual differing from the current solution candidate in only one property. If the new individual is better, it replaces its ancestor, otherwise it is discarded.

¹¹ http://en.wikipedia.org/wiki/A%2A_search [accessed 2007-08-09]

¹² see Definition 26.28 on page 414

- *Greedy dynamics*: The optimization process chooses a single new individual from the set of “one-mutant change” neighbors. If it is not better than the current solution candidate, the search continues until a better one has been found or all neighbors have been enumerated. The major difference to the previous form is the number of steps that are needed per improvement.
- *Fitter Dynamics*: The optimization process enumerates all one-mutant neighbors of the current solution candidate and transcends to the best one.

From these elaborations, it becomes clear that adaptive walks are very similar to hill climbing and Random Optimization. The major difference is that an adaptive walk is a theoretical construct that, very much like random walks, helps us to determine properties of fitness landscapes whereas the other two are practical realizations of optimization algorithms.

Adaptive walks are a very common construct in evolutionary biology. Biological populations are running for a very long time and so their genetic compositions are assumed to be relatively converged [1457, 186]. The dynamics of such populations in near-equilibrium states with low mutation rates can be approximated with one-mutant adaptive walks [1458, 1457, 186].

Parallelization and Distribution

As already stated many times, global optimization problems are often computational intense. Up until now, we have only explored the structure and functionality of optimization algorithms without paying attention to their potential of parallelization or even distribution¹.

Roughly speaking, parallelization² means to search for pieces of code that can potentially run concurrently and letting them execute by different processors [1459, 1460]. Take painting a fence for example. Here, the overall progress will be much faster if more than one painter applies the color to the wood. Distribution³ is a special case of parallelization where the different processors are located on different machines in a network [1461, 1462]. Imagine that each fence-painter would take a piece of the fence to his workshop where he can use a special airbrush which can color the whole piece at once. Distribution comes with the trade-off of additional communication costs for transporting the data, but has the benefit that it is more generic. At the current time, off-the-shelf PCs usually have not more than two CPUs. This limits the benefit of local parallelization. We can, however, connect arbitrarily many of such computers in a network for distributed processing.

17.1 Analysis

In order to understand which parts of an optimization algorithm can be parallelized, the first step is an analysis. We will do such an analysis for evolutionary algorithms as example for population-based optimizers.⁴ The parallelization and distribution of evolutionary algorithms has long been a subject to study and has been discussed by multiple researchers like Alba and Tomassini [1463], Cant'u-Paz [1464; 759], Tan et al. [1465], Tanese [1466], Mühlenbein [1467], and Bollini and Piastra [1468].

There are two components of evolutionary algorithms whose performance potentially can remarkably be increased by parallelization: the evaluation and the reproduction stages. As sketched in Figure 17.1, evaluation is a per-individual process. The values of the objective functions are determined for each solution candidate independently from the rest of the population. Evaluating the individuals often involves complicated simulations and calculations and is thus usually the most time-consuming part of evolutionary algorithms.

During the fitness assignment process, it is normally required to compare solution candidates with the rest of the population, to compute special sets of individuals, or to update some data structures. This makes it very hard for parallelization to provide any speedup.

¹ Section 29.2 on page 485 gives a detailed introduction into distributed algorithms, their advantages and drawbacks.

² <http://en.wikipedia.org/wiki/Parallelization> [accessed 2007-07-03]

³ http://en.wikipedia.org/wiki/Distributed_computing [accessed 2007-11-30]

⁴ In Section 2.1.1 on page 71 you can find the basic evolutionary algorithm.

The selection phase may or may not require access to certain subsets of the population or data updates. Whether parallelization is possible or is beneficial thus depends on the selection scheme applied.

The reproduction phase, on the other hand, can very easily be parallelized. It involves creating a new individual by using (but not altering) the information from n existing ones, where $n = 0$ corresponds to the creation operation, $n = 1$ resembles mutation, and $n = 2$ means recombination. Thus, the making of each new genotype is an independent task.




Fig. 17.1: Parallelization potential in evolutionary algorithm.

Despite running an evolutionary algorithm in single a thread⁵ of execution (see Figure 17.2), our analysis has shown that makes sense to have at least the evaluation and reproduction phase executed in parallel as illustrated in Figure 17.3. Usually, the population is larger than the number of available CPUs⁶, so one thread could be created per processors that consecutively pulls individuals out of a queue and processes them. This approach even yields performance gains on off-the-shelf personal computers since these nowadays at least come with hyper-threading⁷ technology [1469, 1470] or even dual-core⁸ CPUs [1471, 1472].




Fig. 17.2: A sequentially proceeding evolutionary algorithm.

⁵ http://en.wikipedia.org/wiki/Thread_%28computer_science%29 [accessed 2007-07-03]

⁶ <http://en.wikipedia.org/wiki/Cpu> [accessed 2007-07-03]

⁷ <http://en.wikipedia.org/wiki/Hyper-threading> [accessed 2007-07-03]

⁸ <http://en.wikipedia.org/wiki/Dual-core> [accessed 2007-07-03]




Fig. 17.3: A parallel evolutionary algorithm with two worker threads.

Cant'u-Paz [1473] divides parallel evolutionary algorithms into two main classes:

- In *globally* parallelized EAs, each individual in the population can (possibly) always mate with any other.
- In *coarse grained* approaches, the population is divided into several sub-populations where mating inside a sub-population is unrestricted but mating between individuals of different sub-populations may only take place occasionally according to some rule.

In ancient Greece, a *deme* was a district or township inhabited by a group that formed an independent community. They were the basic units of government in Attica as remodeled by Cleisthenes around 500 BC. In biology, a deme is a locally interbreeding group within a geographic population.

Definition 17.1 (Deme). In evolutionary algorithms, a *deme* is a distinct sub-population.

In the following, we are going to discuss some of the different parallelization methods from the viewpoint of distribution because of its greater generality.

17.2 Distribution

The distribution of an algorithm only pays off if the delay induced by the transmissions necessary for data exchange is much smaller than the time saved by distributing the computational load. Thus, in some cases distributing of optimization is useless. If searching for the root of a mathematical function for example, transmitting the parameter vector x to another computer will take much longer than computing the function $f(x)$ locally. In this section, we will investigate some basic means to distribute evolutionary algorithms that can as well as be applied to other optimization methods as outlined by Weise and Geihs [862].

17.2.1 Client-Server

If the evaluation of the objective functions is time consuming, the easiest approach to distribute an evolutionary algorithm is the client-server scheme (also called master-slave).⁹ Figure 17.4 illustrates how we can make use of this very basic, global distribution scheme.

⁹ A general discussion concerning the client-server architecture can be found in Section 29.2.2 on page 488

Here, the servers (slaves) receive the single tasks, process them, and return the results. Such a task can, for example, be the reproduction of one or two individuals and the subsequent determination of the objective values of the offspring. The client (or master) just needs to distribute the parent individuals to the servers and receives their fully evaluated offspring in return. These offspring are then integrated into the population, where fitness assignment and selection is performed. Client-server-based distribution approaches for evolutionary algorithms have been discussed and tested by Van Veldhuizen et al. [1474], Xu et al. [1475], Dubreuil et al. [1476] and were realized in general-purpose software packages by Cahon et al. [1477], Luke et al. [1012], and Weise and Geihs [862].




Fig. 17.4: An EA distributed according to the client-server approach.

One practical realization of this approach can be to use a queue where all the selected individuals are pushed into as mating pool. Each server in the network is then represented by a thread on the client side. Such a thread pulls individuals from the queue, sends them to its corresponding server, and waits for the result to be returned. It places the individuals it receives into the new population and then starts over again. Servers may possess multiple processors, which can be taken into account by representing them by an appropriate number of threads.

17.2.2 Island Model

Under some circumstances, the client-server approach may not be optimal, especially if

- Processing of the tasks is fast relatively to the amount of time needed for the data exchange between the client and the server. In other words, if messages that have to be exchanged travel longer than the work would take if performed locally, the client-server method would actually slow down the system.
- Populations are required that cannot be held completely in the memory of a single computer. This can be the case either if the solution candidates are complex and memory consuming or the nature of the problem requires large populations.

In such cases, we can again learn from nature. Until now we only have imitated evolution on one large continent. All individuals in the population compete with each other and there are no barriers between the solution candidates. In reality, there occur obstacles like mountain ranges or oceans, separating parts of the population and creating isolated sub-populations. Another example for such a scenario is an archipelago like the Galapagos islands where Darwin [6], the father of the evolution theory, performed his studies. On the single islands, different species can evolve independently. From time to time, a few individuals from one isle *migrate* another one, maybe by traveling on a tree trunk over the water or

by been blown there by a storm. If they are fit enough in their new environment, they can compete with the local population and survive. Otherwise, they will be extruded by the native residents of the habitat. This way, the islands manage an approximately equal level of fitness of their individuals, while still preserving a large amount of diversity.




Fig. 17.5: An evolutionary algorithm distributed in a P2P network.

We can easily copy this natural role model in evolutionary algorithms by using multiple sub-populations (demes) as discussed by Cohoon et al. [159], [160], Skolicki and De Jong [1478], [1479], Gorges-Schleuter [1480], Tanese [1466], and Toshine et al. [1481] and also realized in various software packages such as those created by Whitley and Starkweather [1482], Paechter et al. [1483], Tan et al. [1465], Arenas et al. [1484], Chong and Langdon [1485], Luke et al. [1012], Cahon et al. [1477], and Weise and Geihs [862]. Distributing the demes on n different nodes in a network of computers, each representing one island, is maybe the most popular form of coarse grained parallelization. Hence, both disadvantages of the original master/slave approach are circumvented: Communication between nodes is only needed when individuals *migrate* between them. This communication can be performed asynchronously to the n independently running evolutionary algorithms and does not slow down their performance. The migration rule can furthermore be chosen in a way that reduces the network traffic. By dividing the population, the number of solution candidates to be held on single machines also decreases, which helps to mitigate the memory consumption problem.

The island model can be realized by peer-to-peer networks¹⁰ where each node runs an independent evolution, as illustrated in Figure 17.5. Here, we have modified the selection phase which now returns some additional individuals to be transmitted to another node in the system. Depending on the optimization problems, solution candidates migrating over the network can either enter the fitness assignment process on the receiving machine directly or may take part in the evaluation process first. If the latter is the case, different objective functions can be applied on different nodes.

¹⁰ P2P networks are discussed in Section 29.2.2 on page 489.

Driving this thought further, one will recognize that the peer-to-peer approach inherently allows mixing of different optimization technologies, as outlined by Weise and Geihs [862]. On one node, for instance, the SPEA2-algorithm (see ?? on page ??) can be performed, whereas another node could optimize according to plain hill climbing as described in Chapter 8 on page 209. Such a system, illustrated in Figure 17.6, has a striking advantage. From the No Free Lunch Theorem discussed in Section 1.4.9 on page 54, we know that for optimization algorithms perform differently for different problems. If the problem is unimodal, i. e., has exactly one global optimum and no local optima, a hill climbing approach will outperform any other technique since it directly converges to this optimum. If the fitness landscape is rugged, one the other hand, methods like SPEA2 which have a very balanced exploration/exploitation proportion are able to yield better results and hill climbing algorithms will get stuck to local optima. In most cases, it is not possible to know beforehand which optimization strategy will perform best. Furthermore, the best approach may even change while the optimization proceeds. If a new, better individual evolves, i. e., a new optimum is approached, hill climbing will be fast in developing this solution candidate further until its best form is found, i. e., the bottom of the local optimum is reached. In other phases, an exploration of the solution space may be required since all known local optima have been tracked. A technology like Ant Colony Optimization could now come into play. A heterogeneous mixture of these algorithms that exchanges individuals from time to time will retain the good properties of the single algorithms and, in many cases, outperform a homogeneous search [677, 1486, 1487, 1488, 1489]. Just remember how our discussion of Memetic Algorithms in Chapter 15 on page 241.




Fig. 17.6: An example for a heterogeneous search.

The island model can also be applied locally by simply using disjoint local populations. Although this would not bring a performance gain, it could improve the convergence behavior of the optimization algorithm. Spieth et al. [92], for instance, argue that the island model can be used to preserve the solution diversity. By doing so it decreases the probability of premature convergence (see Section 1.4.1 on page 37).

Broadcast-Distributed Parallel Evolutionary Algorithm

The Broadcast-Distributed Parallel Evolutionary Algorithm (BDP) defined by Johnson et al. [1490] extends the island model for large networks of wirelessly connected, resource-restricted devices such as sensor networks, amorphous and paintable computing systems. In the BDP, each node carries a separate population from which one individual is selected after each generation. This individual is broadcasted to the neighbors of the node. Every time a node

receives an individual, it appends it to an internal mate list. Whenever the length of this list exceeds a certain threshold, selection and subsequent crossover is performed on the joint set of the population and the individuals in the mate list.

17.2.3 Mixed Distribution

Of course, we can combine the both distribution approaches previously discussed by having a peer-to-peer network that also contains client-server systems, as sketched in Figure 17.7. Such a system will be especially powerful if we need large populations of individuals that take long to evaluate. Then, the single nodes in the peer-to-peer network together provide a larger virtual population, while speeding up their local evolutions by distributing the computational load to multiple servers.




Fig. 17.7: A mixed distributed evolutionary algorithms.

17.3 Cellular Genetic Algorithms

Cellular Genetic Algorithms are a special family of parallelization models for genetic algorithms which has been studied by various researchers such as Whitley [1491; 645], Manderick and Spiessens [1492], Hillis [1493], and Davidor [1494]. A good understanding of this model can be reached by starting with the basic architecture of the cellular system as described by Whitley [645].

Assume we have a matrix m of $N \times N$ cells. Each cell has a processor and holds one individual of the population. It can communicate with its right, left, top, and bottom neighbor. Cells at the edge of the matrix are wired with cells in the same column/row at the

opposite edge. The cell $m_{I,j}$ can thus communicate with $m_{(i+1) \bmod N,j}$, $m_{(i-1) \bmod N,j}$, $m_{i,(j+1) \bmod N}$, and $m_{i,(j-1) \bmod N}$. It is also possible to extend this neighborhood to all cells in a given Manhattan distance, but let us stick with the easiest case.

Each cell can evaluate the individual it locally holds. For creating offspring, it can either mutate this individual or recombine it with one selected from of the four solution candidates on its neighbors. At the beginning, the matrix is initialized with random individuals. After some time, the spatial restriction of mating leads to the occurrence of local neighborhoods with similar solution candidates denoting local optima. These hoods begin to grow until they touch each other. Then, the regions better optima will “consume” worse ones and reduce the overall diversity.

Although there are no fixed mating restrictions like in the island model, regions that are about twenty or so moves away will virtually not influence each other. We can consider groups of cells that distant as separate sub-populations. This form of separation is called *isolation by distance* – again, a term that originally stems from biology (coined by Wright [48]). [645, 1495, 1467, 1480] For observing such effects, it is said that a certain minimum of cells is required – at least about 1000 according to Whitley [645].

Maintaining the Optimal Set

Most multi-objective optimization algorithms return a set of optimal solutions X^* instead of a single individual x^* . Many optimization techniques also internally keep track of the set of best solution candidates encountered during the search process. In Simulated Annealing, for instance, it is quite possible to discover an optimal element $\mathbf{x}^*$ and subsequently depart from it to a local optimum x_l^* . Therefore, optimizers normally carry a list of the non-prevailed solution candidates ever visited with them.

In scenarios where the search space $\mathbb{G}$ differs from the problem space $\mathbb{X}$ it often makes more sense to store the list of optimal individual records P^* instead of just keeping the optimal phenotypes x^* . Since the elements of the search space are no longer required at the end of the optimization process, we define the simple operation `extractPhenotypes` which extracts them from a set of individuals P .

$$\forall x \in \text{extractPhenotypes}(P) \Rightarrow \exists p \in P : x = p.x \quad (18.1)$$

18.1 Updating the Optimal Set

Whenever a new individual p is created, the set of optimal individuals P^* may change. It is possible that the new solution candidate must be included in the optimal set or even prevails some of the phenotypes already contained therein which then must be removed.

Definition 18.1 (`updateOptimalSet`). The function `updateOptimalSet` updates a set of optimal elements P_{old}^* with the new solution candidate $p_{new}.x$. It uses implicit knowledge of the prevalence relation $\succ$ and the corresponding comparator function `cmpF`.

$$\begin{aligned} P_{new}^* &= \text{updateOptimalSet}(P_{old}^*, p_{new}), \\ P_{old}^*, P_{new}^* &\subseteq \mathbb{G} \times \mathbb{X}, p_{new} \in \mathbb{G} \times \mathbb{X} : \\ \forall p_1 \in P_{old}^* \nexists p_2 \in P_{old}^* : p_2.x \succ p_1.x &\Rightarrow P_{new}^* \subseteq P_{old}^* \cup \{p_{new}\} \wedge \\ &\forall p_1 \in P_{new}^* \nexists p_2 \in P_{new}^* : p_2.x \succ p_1.x \end{aligned} \quad (18.2)$$

We define two equivalent approaches in Algorithm 18.1 and Algorithm 18.2 which perform the necessary operations. Algorithm 18.1 creates a new, empty optimal set and successively inserts optimal elements whereas Algorithm 18.2 removes all elements which are prevailed by the new individual p_{new} from the old optimal set P_{old}^* .

Especially in the case of evolutionary algorithms, not a single new element is created in each generation but a set P . Let us define the operation `updateOptimalSetN` for this purpose. This operation can easily be realized by iteratively applying `updateOptimalSet`, as shown in Algorithm 18.3.

Algorithm 18.1: $P_{new}^* \leftarrow \text{updateOptimalSet}(P_{old}^*, p_{new})$

Input: P_{old}^* : the optimal set as known before the creation of p_{new}
Input: p_{new} : a new individual to be checked
Output: P_{new}^* : the optimal set updated with the knowledge of p_{new}

```

1 begin
2    $P_{new}^* \leftarrow \emptyset$ 
3   foreach  $p_{old} \in P_{old}^*$  do
4     if  $p_{new}.x \succ p_{old}.x$  then
5        $P_{new}^* \leftarrow P_{new}^* \cup \{p_{old}\}$ 
6     if  $p_{old}.x \succ p_{new}.x$  then return  $P_{old}^*$ 
7   return  $P_{new}^* \cup \{p_{new}\}$ 
8 end

```

Algorithm 18.2: $P_{new}^* \leftarrow \text{updateOptimalSet}(P_{old}^*, p_{new})$ (*2nd Version*)

Input: P_{old}^* : the optimal set as known before the creation of p_{new}
Input: p_{new} : a new individual to be checked
Output: P_{new}^* : the optimal set updated with the knowledge of p_{new}

```

1 begin
2    $P_{new}^* \leftarrow P_{old}^*$ 
3   foreach  $p_{old} \in P_{old}^*$  do
4     if  $p_{new}.x \succ p_{old}.x$  then
5        $P_{new}^* \leftarrow P_{new}^* \setminus \{p_{old}\}$ 
6     else if  $p_{old}.x \succ p_{new}.x$  then
7       return  $P_{old}^*$ 
8   return  $P_{new}^* \cup \{p_{new}\}$ 
9 end

```

Algorithm 18.3: $P_{new}^* \leftarrow \text{updateOptimalSetN}(P_{old}^*, P)$

Input: P_{old}^* : the old optimal set
Input: P : the set of new individuals to be checked for optimality
Data: p : an individual from P
Output: P_{new}^* : the updated optimal set

```

1 begin
2    $P_{new}^* \leftarrow P_{old}^*$ 
3   foreach  $p \in P$  do  $P_{new}^* \leftarrow \text{updateOptimalSet}(P_{new}^*, p)$ 
4   return  $P_{new}^*$ 
5 end

```

18.2 Obtaining Optimal Elements

The function `updateOptimalSet` helps an optimizer to build and maintain a list of optimal individuals. When the optimization process finishes, the `extractPhenotypes` can then be used to obtain the optimal elements of the problem space and return them to the user. However, not all optimization methods maintain an optimal set all the time. When they terminate, they have to extract all optimal elements the set of individuals Pop currently known.

Definition 18.2 (`extractOptimalSet`). The function `extractOptimalSet` function extracts a set P^* of optimal (non-prevalled) individuals from any given set of individuals Pop .

$$\forall P^* \subseteq Pop \subseteq \mathbb{G} \times \mathbb{X}, P^* = \text{extractOptimalSet}(Pop) \Rightarrow \forall p_1 \in P^* \nexists p_2 \in Pop : p_2.x \succ p_1.x \quad (18.3)$$

Algorithm 18.4 defines one possible realization of `extractOptimalSet`. By the way, this approach could also be used for updating an optimal set.

$$\text{updateOptimalSet}(P_{old}^*, p_{new}) \equiv \text{extractOptimalSet}(P_{old}^* \cup p_{new}) \quad (18.4)$$

Algorithm 18.4: $P^* \leftarrow \text{extractOptimalSet}(Pop)$

Input: Pop : the list to extract the optimal individuals from
Data: p_{any}, p_{chk} : solution candidates tested for supremacy
Data: i, j : counter variables
Output: P^* : the optimal subset extracted from Pop

```

1 begin
2    $P^* \leftarrow Pop$ 
3   for  $i \leftarrow \text{len}(P^*) - 1$  down to 1 do
4     for  $j \leftarrow i - 1$  down to 0 do
5       if  $P^*[i] \succ P^*[j]$  then
6          $P^* \leftarrow \text{deleteListItem}(P^*, j)$ 
7          $i \leftarrow i - 1$ 
8       else if  $P^*[j] \succ P^*[i]$  then
9          $P^* \leftarrow \text{deleteListItem}(P^*, i)$ 
10         $j \leftarrow -1$ 
11   return listToSet( $P^*$ )
12 end
```

18.3 Pruning the Optimal Set

In some optimization problems, there may be very many if not infinite many optimal individuals. The set of X^* optimal solution candidates computed by the optimization algorithms, however, cannot grow infinitely because we only have limited memory. Therefore, we need to perform an action called *pruning* which reduces the size of the optimal set to a given limit k [1496, 1497, 1498].

There exists a variety of possible pruning operations. Morse [1496] and Taboada and Coit [1499], for instance, suggest to use clustering algorithms¹ for this purpose. In principle, also any combination of the fitness assignment and selection schemes discussed in Chapter 2 would do. It is very important that the loss of generality during a pruning operation is minimized. Fieldsend et al. [1500], for instance, point out that if the extreme elements of the optimal frontier are lost, the resulting set may only represent a small fraction of the optimal that could have been found without pruning. Instead of working on the set of optimal solution candidates X^* , we again base our approach on a set of optimal individuals P^* and we define:

Definition 18.3 (`pruneOptimalSet`). The pruning operation `pruneOptimalSet` reduces the size of a set P_{old}^* of individuals to fit to a given upper boundary k .

$$\forall P_{new}^* \subseteq P_{old}^* \subseteq \mathbb{G} \times \mathbb{X}, k \in \mathbb{N} : P_{new}^* = \text{pruneOptimalSet}(P_{old}^*, k) \Rightarrow |P_{new}^*| \leq k \quad (18.5)$$

¹ You can find a discussion of clustering algorithms in Section 28.2.

18.3.1 Pruning via Clustering

Algorithm 18.5 uses clustering to provide the functionality specified in this definition and thereby realizes the idea of Morse [1496] and Taboada and Coit [1499]. Basically, any given clustering algorithm could be used as replacement for cluster – see Chapter 28 on page 467 for more information on clustering.

Algorithm 18.5: $P_{new}^* \leftarrow \text{pruneOptimalSet}_c(P_{old}^*, k)$

Input: P_{old}^* : the optimal set to be pruned
Input: k : the maximum size allowed for the optimal set ($k > 0$)
Input: [implicit] cluster: the clustering algorithm to be used
Input: [implicit] nucleus: the function used to determine the nuclei of the clusters
Data: B : the set of clusters obtained by the clustering algorithm
Data: b : a single cluster $b \in B$
Output: P_{new}^* : the pruned optimal set

```

1 begin
    // obtain  $k$  clusters
2    $B \leftarrow \text{cluster}(P_{old}^*)$ 
3    $P_{new}^* \leftarrow \emptyset$ 
4   foreach  $b \in B$  do  $P_{new}^* \leftarrow P_{new}^* \cup \text{nucleus}(b)$ 
5   return  $P_{new}^*$ 
6 end

```

18.3.2 Adaptive Grid Archiving

Let us discuss the adaptive grid archiving algorithm (AGA) as example for a more sophisticated approach to prune the optimal set. AGA has been introduced for the evolutionary algorithm PAES² by Knowles and Corne [1501] and uses the objective values (computed by the set of objective functions F) directly. Hence, it can treat the individuals as $|F|$ -dimensional vectors where each dimension corresponds to one objective function $f \in F$. This $|F|$ -dimensional objective space $\mathbb{Y}$ is divided in a grid with d divisions in each dimension. Its span in each dimension is defined by the corresponding minimum and maximum objective values. The individuals with the minimum/maximum values are always preserved. This circumvents the phenomenon of narrowing down the optimal set described by Fieldsend et al. [1500] and distinguishes the AGA approach from clustering-based methods. Hence, it is not possible to define maximum optimal set sizes k which are smaller than $2|F|$. If individuals need to be removed from the set because it became too large, the AGA approach removes those that reside in regions which are the most crowded.

The original sources outline the algorithm basically with with descriptions and definitions. Here, we introduce a more or less trivial specification in Algorithm 18.6 on the facing page and Algorithm 18.7 on page 268. The function `agaDivide` is internally used to perform the grid division. It transforms the objective values of each individual to grid coordinates stored in the array *lst*. Furthermore, `agaDivide` also counts the number of individuals that reside in the same coordinates for each individual and makes it available in *cnt*. It ensures the preservation of border individuals by assigning a negative *cnt* value to them. This basically disables their later disposal by the pruning algorithm `pruneOptimalSetaga` since `pruneOptimalSetaga` deletes the individuals from the set P_{old}^* that have largest *cnt* values first.

² PAES is discussed in ?? on page ??

Algorithm 18.6: $(P_l, lst, cnt) \leftarrow \text{agaDivide}(P_{old}, d)$

Input: P_{old} : the optimal set to be pruned**Input:** d : the number of divisions to be performed per dimension**Input:** $[implicit] F$: the set of objective functions**Data:** i, j : counter variables**Data:** $mini, maxi, mul$: temporary stores**Output:** (P_l, lst, cnt) : a tuple containing the list representation P_l of P_{old} , a list lst assigning grid coordinates to the elements of P_l and a list cnt containing the number of elements in those grid locations

```

1 begin
2    $mini \leftarrow \text{createList}(|F|, \infty)$ 
3    $maxi \leftarrow \text{createList}(|F|, -\infty)$ 
4   for  $i \leftarrow |F| - 1$  down to 0 do
5      $mini[i-1] \leftarrow \min \{f_i(p.x) \mid \forall p \in P_{old}\}$ 
6      $maxi[i-1] \leftarrow \max \{f_i(p.x) \mid \forall p \in P_{old}\}$ 
7    $mul \leftarrow \text{createList}(|F|, 0)$ 
8   for  $i \leftarrow |F| - 1$  down to 0 do
9     if  $maxi[i] \neq mini[i]$  then
10       $mul[i] \leftarrow \frac{d}{maxi[i] - mini[i]}$ 
11    else
12       $maxi[i] \leftarrow maxi[i] + 1$ 
13       $mini[i] \leftarrow mini[i] - 1$ 
14    $P_l \leftarrow \text{setToList}(P_{old})$ 
15    $lst \leftarrow \text{createList}(\text{len}(P_l), \emptyset)$ 
16    $cnt \leftarrow \text{createList}(\text{len}(P_l), 0)$ 
17   for  $i \leftarrow \text{len}(P_l) - 1$  down to 0 do
18      $lst[i] \leftarrow \text{createList}(|F|, 0)$ 
19     for  $j \leftarrow 1$  up to  $|F|$  do
20       if  $(f_j(P_l[i]) \leq mini[j-1]) \vee (f_j(P_l[i]) \geq maxi[j-1])$  then
21          $cnt[i] \leftarrow cnt[i] - 2$ 
22          $lst[i][j-1] \leftarrow \lfloor (f_j(P_l[i]) - mini[j-1]) * mul[j-1] \rfloor$ 
23       if  $cnt[i] > 0$  then
24         for  $j \leftarrow i + 1$  up to  $\text{len}(P_l) - 1$  do
25           if  $lst[i] = lst[j]$  then
26              $cnt[i] \leftarrow cnt[i] + 1$ 
27             if  $cnt[j] > 0$  then  $cnt[j] \leftarrow cnt[j] + 1$ 
28   return  $(P_l, lst, cnt)$ 
29 end

```

Algorithm 18.7: $P_{new}^* \leftarrow \text{pruneOptimalSet}_{aga}(P_{old}^*, d, k)$

Input: P_{old}^* : the optimal set to be pruned**Input:** d : the number of divisions to be performed per dimension**Input:** k : the maximum size allowed for the optimal set ($k \geq 2|F|$)**Input:** F : the set of objective functions**Data:** i : a counter variable**Data:** P_l : the list representation of P_{old}^* **Data:** lst : a list assigning grid coordinates to the elements of P_l **Data:** cnt : the number of elements in the grid locations defined in lst **Output:** P_{new}^* : the pruned optimal set

```

1 begin
2   if  $\text{len}(P_{old}^*) \leq k$  then return  $P_{old}^*$ 
3    $(P_l, lst, cnt) \leftarrow \text{agaDivide}(P_{old}^*, d)$ 
4   while  $\text{len}(P_l) > k$  do
5      $idx \leftarrow 0$ 
6     for  $i \leftarrow \text{len}(P_l) - 1$  down to 1 do
7       if  $cnt[i] > cnt[idx]$  then  $idx \leftarrow i$ 
8     for  $i \leftarrow \text{len}(P_l) - 1$  down to 0 do
9       if  $(lst[i] = lst[idx]) \wedge (cnt[i] > 0)$  then  $cnt[i] \leftarrow cnt[i] - 1$ 
10     $P_l \leftarrow \text{deleteListItem}(P_l, idx)$ 
11     $cnt \leftarrow \text{deleteListItem}(cnt, idx)$ 
12     $lst \leftarrow \text{deleteListItem}(lst, idx)$ 
13  return  $\text{listToSet}(P_l)$ 
14 end

```

Part II

Applications

Experimental Settings, Measures, and Evaluations

In this chapter we will discuss the possible experimental settings, the things that we can measure during experiments, and what information we can extract from these measurements. We will also define some suitable shortcuts which we use later in descriptions of experiments in order to save space in tables and graphics.

19.1 Settings

Experiments can only be reproduced and understood if their setup is precisely described. Especially in the area of global optimization algorithms, there is a wide variety of possible parameter settings. Incompletely documenting an experiment may lead to misunderstandings. Other researchers repeating the tests will use default settings for all parameters not specified (or any settings that they find neat) and possibly obtain totally different results. Here, we will try to list many possible parameters of experiments (without claiming completeness). Of course, not all of them are relevant in a given experiment. Instead, this list is to be understood as a hint of what to consider when specifying a test series. In many tables and graphics, we will use shortcuts of the parameter names in order to save space.

19.1.1 The Optimization Problem

As stated in Definition 1.34 on page 27, an optimization problem is a five-tuple $(\mathbb{X}, F, \mathbb{G}, Op, gpm)$ specifying the problem space $\mathbb{X}$, the objective functions F , the search space $\mathbb{G}$, the set of search operations Op , and the genotype-phenotype mapping gpm . Specifying the elements of this tuple is the most important prerequisite for any experiment. Table 19.1 gives an example for this structure.

Parameter	Short	Description
Problem Space	$\mathbb{X}$	The space of possible solution candidates. (see Section 1.3.1) <i>Example:</i> The variable length natural vectors $x = (x_0, x_1, \dots)^T$, $x_i \in \mathbb{N}_0 \forall i \in [0, \text{len}(x) - 1]$, $0 < \text{len}(x) \leq 500 \forall x \in \mathbb{X}$
Objective Functions	F	The objective functions which measure the utility of the solution candidates. If nothing else is stated, <i>minimization</i> is assumed. (see Definition 1.1) <i>Example:</i> $F = \{f_1, f_2\} : f_1(x) = \sum_{i=0}^{\text{len}(x)-1} x_i$, $f_2(x) = \sum_{i=0}^{\text{len}(x)-1} \sin x_i$

Search Space	$\mathbb{G}$	The space of the elements where the search operations are applied on. (see Section 1.3.1) <i>Example:</i> The variable length bit strings $g : 0 < \text{len}(g) \leq 2000$.
Search Operations	Op	The search operations available for the optimizer. Here it is also important to note the way in which they are applied. (see Section 1.3.1) <i>Example:</i> creation: uniform in length and values $os = 1 \rightarrow 10\% \text{ mutation, } 90\% \text{ multi-point crossover}$ $os = 2 \rightarrow 25\% \text{ mutation, } 80\% \text{ multi-point crossover}$ $os = 3 \rightarrow 45\% \text{ mutation, } 65\% \text{ multi-point crossover}$ The results from crossover may be mutated (non-exclusive search operations).
GPM	gpm	The genotype-phenotype mapping translates points from the search space $\mathbb{G}$ to points in the problem space $\mathbb{X}$. (see Definition 1.30) <i>Example:</i> $x = gpm(g) \Rightarrow x_i = \sum_{j=4i}^{4i+3} g_j, \text{len}(x) = \lfloor \frac{1}{4} \text{len}(g) \rfloor$

Table 19.1: The basic optimization problem settings.

In this table we have defined a simple example optimization problem which has a search space composed of vectors with between 1 and 500 elements (natural numbers). These vectors are encoded as variable length bit strings, where groups of four bits stand for one vector element. As objective functions, two simple sums over the vector elements are applied. When needed, new strings with uniformly distributed length and contents are created with the creation operation. Sometimes, a test series involves tests with different settings. This is the case in this example too, where three configurations for the reproduction operations are given. In a table containing the results of the experiments, there could be a column “ os ” may contain the values from 1 to 3 corresponding to these settings. In our example, elements resulting from the crossover may be mutated (which is meant by “non-exclusive”). Therefore, the percentages in which the operators are applied do not necessarily need to sum up to 100%. With this definition, the problem can be reproduced easily and it is also possible to apply different global optimization algorithms and to compare the results.

19.1.2 The Optimization Algorithm Applied

The performance of an optimization algorithm strongly depends on its configuration. Table 19.1 lists some of the parameters most commonly involved in this book and gives examples how they could be configured.

Parameter	Short	Description
Optimization Algorithm	alg	The optimization algorithm used to solve the problem. (see Definition 1.39) <i>Example:</i> $alg = 0 \rightarrow$ (Parallel) Random Walks $alg = 1 \rightarrow$ evolutionary algorithm $alg = 2 \rightarrow$ Memetic Algorithm $alg = 3 \rightarrow$ Simulated Annealing $alg = 4 \rightarrow$ downhill simplex

Comparison Operator	cm	In multi-objective optimization, individuals are often compared according to Pareto or prevalence schemes. This parameter states which scheme was used, if any. (see Section 1.2.7) <i>Example:</i> $cm = 0 \rightarrow$ weighted sum $cm = 1 \rightarrow$ Pareto comparison
Steady-State	ss	Are the parent individuals in the population simply discarded (generational) or do they compete with their offspring (steady-state). This parameter is usually only valid in the context of evolutionary algorithms or other population-oriented optimizers. (see Section 2.1.2) <i>Example:</i> $ss = 0 \rightarrow$ generational $ss = 1 \rightarrow$ steady-state
Population Size	ps	The population size (only valid for population-oriented algorithms). (see Definition 2.4) <i>Example:</i> $ps \in \{10, 100, 1000, 10\,000\}$
Elitism	el	Are the best solution candidates preserved (elitism) or not (no elitism)? (see Definition 2.4) <i>Example:</i> $el = 0 \rightarrow$ elitism is used $el = 1 \rightarrow$ no elitism is used
Maximum Archive Size	as	The size of the archive with the best known individuals (only valid if elitism is used). Notice: An archive size of zero corresponds to no elitism. (see Definition 2.4) <i>Example:</i> $as \in \{0, 10, 20, 40, 80\}$
Fitness Assignment Algorithm	fa	The fitness assignment algorithm used. (see Section 2.3) <i>Example:</i> $fa = 0 \rightarrow$ weighted sum fitness assignment $fa = 1 \rightarrow$ Pareto ranking $fa = 2 \rightarrow$ Variety Preserving Ranking
Selection Algorithm	sel	The selection algorithm used. (see Section 2.4) <i>Example:</i> $sel = 0 \rightarrow$ Fitness proportionate selection $sel = 1 \rightarrow$ Tournament selection $sel = 2 \rightarrow$ Truncation Selection
Tournament Size	k	The number of individuals competing in tournaments in tournament selection (only valid for $sel = 1$). (see Section 2.4.8) <i>Example:</i> $k \in \{2, 3, 4, 5\}$
Convergence Prevention	cp	Is the simple convergence prevention method used? (see Section 2.4.8) <i>Example:</i> $cp \in \{0, 0.1, 0.2, 0.3, 0.4\}$

Table 19.2: Some basic optimization algorithm settings.

Such a table describes a set of experiments if at least one of the parameters has more than one value. If several parameters can be configured differently, the number of experiments multiplies accordingly. Then, (full) *factorial experiments*¹ [1502, 1503, 1504] where all possible parameter combinations are tested separately (multiple times) can be performed.

19.1.3 Other Run Parameters

In ??, some additional parameters describing how the optimization algorithm was applied and how the experiments were carried out.

Parameter	Short	Description
Number of Training Cases	tc	The number of training cases used for evaluating the objective functions. <i>Example:</i> $tc \in \{1, 10, 20\}$

¹ http://en.wikipedia.org/wiki/Factorial_experiment [accessed 2008-08-07]

Training Case Change Policy	ct	The policy according to which the training cases are changed. (see Definition 1.39) <i>Example:</i> $ct = 0 \rightarrow$ The training cases do not change. $ct = 1 \rightarrow$ The training cases change each generation. $ct = 2 \rightarrow$ Training cases change after each evaluation.
Evaluation Limit	$m\tau$	The maximum number of individual evaluations τ that each run is allowed to perform. (see Definition 1.42) <i>Example:</i> $m\tau = 45\,000$
Generation Limit	mt	The maximum number of iterations/generations t that each run is allowed to perform. (see Definition 1.43) <i>Example:</i> $mt = 1000$
Maximum Number of Runs	mr	The (maximum) number of runs to perform. This threshold can be combined with time constraints which may lead to fewer runs being performed. <i>Example:</i> $mr = 40$
Maximum Time per Run	mrT	The (maximum) amount of time granted per run. <i>Example:</i> $mrT = 40h$
Maximum Total Time	mT	The (maximum) total time granted. <i>Example:</i> $mT = 40d\,5h$
System Configuration	Cfg	Especially in cases where time constraints are imposed, the configuration of the system on which the experiments run becomes important. <i>Example:</i> $Cfg = 0 \rightarrow$ one 9 GHz two-core PC, 3 GiB RAM, Windows XP, Java 1.4 $Cfg = 1 \rightarrow$ one 9 GHz two-core PC, 3 GiB RAM, Windows XP, Java 1.6

Table 19.3: Some additional parameters of experiments.

19.2 Measures

In Table 19.4 we list some basic measurements that easily can be taken from each test series of a single configuration. From these basic results, more meaningful metrics can be computed.

Measure	Short	Description
Number of Comp. Runs	$\#r$	The total number of completed runs with the specified configuration. <i>Example:</i> $\#r = 100$
Success Evaluations	$s\tau_i$	The number of individual evaluations τ performed in run i until the first individual with optimal functional objective values occurred. This individual may have non-optimal non-functional objective values or be overfitted. (see Definition 1.42) <i>Example:</i> $s\tau_i = 100 \rightarrow$ 1 st successful individual in evaluation 100 $s\tau_i = \emptyset \rightarrow$ no successful individual was found

Success Generations	st_i	The number of iterations/generations t performed in run i until the first individual with optimal functional objective values occurred. This individual may have non-optimal non-functional objective values or be overfitted. (see Definition 1.43) <i>Example:</i> $st_i = 800 \rightarrow 1^{st}$ successful individual in generation 800 $st_i = \emptyset \rightarrow$ no successful individual was found
Perfection Evaluation	$p\tau_i$	The number of individual evaluations τ performed in run i until a “perfect” individual with optimal functional objective values occurred. Depending on the context, this may be an individual where all objective values are optimal, a non-overfitted individual with optimal functional objectives, or both. Thus, $pt_i \geq st_i$ always holds (see Definition 1.42) <i>Example:</i> $p\tau_i = 100 \rightarrow 1^{st}$ perfect individual in evaluation 100 $p\tau_i = \emptyset \rightarrow$ no perfect individual was found
Perfection Generation	pt_i	The number of iterations/generations t performed in run i until a “perfect” individual with optimal functional objective values occurred. Depending on the context, this may be an individual where all objective values are optimal, a non-overfitted individual with optimal functional objectives, or both. Thus, $pt_i \geq st_i$ always holds (see Definition 1.43) <i>Example:</i> $pt_i = 800 \rightarrow 1^{st}$ perfect individual in generation 800 $pt_i = \emptyset \rightarrow$ no perfect individual was found
Solution Set	X_i^*	The set $X_i^* \subseteq \mathbb{X}$ of solutions returned by run i. <i>Example:</i> $X_i^* = \{(0, 1, 2)^T, (4, 5, 6)^T\}$
Runtime	rT_i	The total time needed by run i. <i>Example:</i> $rT_i = 312s$

Table 19.4: Some basic measures that can be obtained from experiments.

The distinction between successful and perfect individuals may seem confusing at first glance and is not necessary in many experiments. Often though, such a distinction is useful. Assume, for instance, that we want to optimize a schedule for a transportation or manufacturing company. A successful schedule, in this case, would be one that allows the company to process all orders. This does not necessarily mean that it is optimal. Perfect would stand for optimal in this context while successful would translate to feasible. If we evolve programs with Genetic Programming and use training cases for their evaluation, we can consider a run as successful if it finds a program which works correctly on all training cases. This, however, could also be caused by overfitting. Then, a perfect program would be one that either also works correctly on another, larger set of test cases not used during the optimization process or whose correctness has been proven. An experimental run is called “successful” (“perfect”) if it found a successful (“perfect”) individual. The concrete definition of successful and perfect is problem specific and must be stated whenever using these predicates. Furthermore notice that we use the sign $\emptyset$ in order to denote runs where no successful (or perfect) solution candidate was discovered.

19.3 Evaluation

19.3.1 Simple Evaluation Measures

After a series of experiments has been carried out and the simple measures from Section 19.2 have been collected, we can use them to compute some first metrics listed in Table 19.5. Obviously, not all of them are needed or carry a meaning in every experiment.

Measure	Short	Description
Number of Successful Runs	$\#s$	The number of runs where successful individuals were discovered. $\#s = \{i : (s\tau_i \neq \emptyset) \wedge (0 \leq i < \#r)\} \quad (19.1)$
Success Fraction	s/r	The fraction of experimental runs that turned out successful. $s/r = \frac{\#s}{\#r} \quad (19.2)$
Minimum Success Evaluation	$\widetilde{s\tau}$	The number of evaluations τ needed by the fastest (successful) experimental run to find a successful individual. (or $\emptyset$ if no run was successful) $\widetilde{s\tau} = \begin{cases} \min \{s\tau_i \neq \emptyset\} & \text{if } \exists s\tau_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.3)$
Mean Success Evaluation	$\overline{s\tau}$	The average number of evaluations τ needed by the (successful) experimental runs to find a successful individual. (or $\emptyset$ if no run was successful) $\overline{s\tau} = \begin{cases} \frac{\sum_{s\tau_i \neq \emptyset} s\tau_i}{ \{s\tau_i \neq \emptyset\} } & \text{if } \exists s\tau_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.4)$
Maximum Success Evaluation	$\widehat{s\tau}$	The number of evaluations τ needed by the slowest (successful) experimental run to find a successful individual. (or $\emptyset$ if no run was successful) $\widehat{s\tau} = \begin{cases} \max \{s\tau_i \neq \emptyset\} & \text{if } \exists s\tau_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.5)$
Minimum Success Generation	$\widetilde{st}$	The number of generations/iterations t needed by the fastest (successful) experimental run to find a successful individual. (or $\emptyset$ if no run was successful) $\widetilde{st} = \begin{cases} \min \{st_i \neq \emptyset\} & \text{if } \exists st_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.6)$
Mean Success Generation	$\overline{st}$	The average number of generations/iterations t needed by the (successful) experimental runs to find a successful individual. (or $\emptyset$ if no run was successful) $\overline{st} = \begin{cases} \frac{\sum_{st_i \neq \emptyset} st_i}{ \{st_i \neq \emptyset\} } & \text{if } \exists st_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.7)$
Maximum Success Generation	$\widehat{st}$	The number of Generations generations/iterations t needed by the slowest (successful) experimental run to find a successful individual. (or $\emptyset$ if no run was successful) $\widehat{st} = \begin{cases} \max \{st_i \neq \emptyset\} & \text{if } \exists st_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.8)$
Number of Perfect Runs	$\#p$	The number of runs where perfect individuals were discovered. $\#p = \{i : (p\tau_i \neq \emptyset) \wedge (0 \leq i < \#r)\} \quad (19.9)$
Perfection Fraction	p/r	The fraction of experimental runs that found perfect individuals. $p/r = \frac{\#p}{\#r} \quad (19.10)$
Minimum Perfection Evaluation	$\widetilde{p\tau}$	The number of evaluations τ needed by the fastest (perfect) experimental run to find a perfect individual. (or $\emptyset$ if no run was found one) $\widetilde{p\tau} = \begin{cases} \min \{p\tau_i \neq \emptyset\} & \text{if } \exists p\tau_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.11)$

Mean Perfection Evaluation	$\overline{p\tau}$	The average number of evaluations τ needed by the (perfect) experimental runs to find a perfect individual. (or $\emptyset$ if no run was found one) $\overline{p\tau} = \begin{cases} \frac{\sum_{p\tau_i \neq \emptyset} p\tau_i}{ \{p\tau_i \neq \emptyset\} } & \text{if } \exists p\tau_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.12)$
Maximum Perfection Evaluation	$\widehat{p\tau}$	The number of evaluations τ needed by the slowest (perfect) experimental run to find a perfect individual. (or $\emptyset$ if no run found one) $\widehat{p\tau} = \begin{cases} \max \{p\tau_i \neq \emptyset\} & \text{if } \exists p\tau_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.13)$
Minimum Perfection Generation	$\widetilde{pt}$	The number of generations/iterations t needed by the fastest (perfect) experimental run to find a perfect individual. (or $\emptyset$ if no run was found one) $\widetilde{pt} = \begin{cases} \min \{pt_i \neq \emptyset\} & \text{if } \exists pt_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.14)$
Mean Perfection Generation	$\overline{pt}$	The average number of generations/iterations t needed by the (perfect) experimental runs to find a perfect individual. (or $\emptyset$ if no run was found one) $\overline{pt} = \begin{cases} \frac{\sum_{pt_i \neq \emptyset} pt_i}{ \{pt_i \neq \emptyset\} } & \text{if } \exists pt_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.15)$
Maximum Perfection Generation	$\widehat{pt}$	The number of generations/iterations t needed by the slowest (perfect) experimental run to find a perfect individual. (or $\emptyset$ if no run found one) $\widehat{pt} = \begin{cases} \max \{pt_i \neq \emptyset\} & \text{if } \exists pt_i \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad (19.16)$
Mean Runtime	$\overline{rT}$	The arithmetic mean of the runtime consumed by the single runs. $\overline{rT} = \frac{1}{\#r} \sum_{i=0}^{r-1} rT_i \quad (19.17)$

Table 19.5: Simple evaluation results.

19.3.2 Sophisticated Estimates

The average generation of success $\overline{st}$ is an estimate of the expected number of iterations needed by the optimizer to find a solution to the optimization problem specified. From the measurements taken during the experiments, however, we can also extract some more sophisticated estimates which give us more information about the optimization process.

Cumulative Probability of Success etc.

One of these estimate is the *cumulative probability of success* $CPs(ps, t')$ introduced by Koza [65]. It approximates the probability that a population-oriented optimization algorithm with a population size ps solves a given problem until iteration (generation) t' . Basically, it can easily be estimated from the experimental data as follows:

$$CPs(ps, t') = \frac{|\{st_i : st_i \leq t'\}|}{\#r} \quad (19.18)$$

The probability of solving the problem until the t'^{th} iteration at least once in st_n independent runs then becomes approximately $1 - (1 - CPs(ps, t'))^{st_n}$. If we want to find out how many runs with t' iterations we need to solve the problem with probability z , we have to solve $z = 1 - (1 - CPs(ps, t'))^{st_n}$. This equation can be solved and we obtain the function $st_n(z, ps, t')$:

$$st_n(z, ps, t') = \begin{cases} \frac{\log(1-z)}{\log(1-CPs(ps, t'))} & \text{if } 0 < CPs(ps, t') < z \\ 1 & \text{if } CPs(ps, t') \geq z \\ +\infty & \text{otherwise} \end{cases} \quad (19.19)$$

From this value, we can directly compute an estimate of the number of objective function evaluations $st_n(z, ps, t')$ needed (i. e., the individuals to be processed) to find a solution with probability z if $st_n(z, ps, t')$ independent runs proceed up to t' iterations. If we have a generational evolutionary algorithm (i. e., $ss = 1$), this would be $st_n(z, ps, t') * ps * t'$. In a steady-state algorithm where all ps parents compete with a total of ps offspring and tc training cases are used that change each generation ($ct = 1$), we would require $2 * ps * tc * st_n(z, ps, t') * t'$ evaluations. If the training cases were constant, we would not need to re-evaluate the parents in order to determine their objective values, and $st_n(z, ps, t')$ would become $ps * tc * st_n(z, ps, t') * t'$.

Obviously, how the value of $st_n(z, ps, t')$ is computed strongly depends on the optimization algorithm applied and may be different from case to case. For the sake of simplicity, we assume that even if we have $|F| > 1$ objective functions, all of them can be evaluated in one single evaluation step together if not explicitly stated otherwise.

Since we often distinguish successful and perfect solutions in this book, we can easily derive the estimates CPp , pt_n , and $p\tau_n$ analogously to CPs , st_n , and $s\tau_n$:

$$CPp(ps, t') = \frac{|\{pt_i : pt_i \leq t'\}|}{\#T} \quad (19.20)$$

$$pt_n(z, ps, t') = \begin{cases} \frac{\log(1-z)}{\log(1-CPp(ps, t'))} & \text{if } 0 < CPp(ps, t') < z \\ 1 & \text{if } CPp(ps, t') \geq z \\ +\infty & \text{otherwise} \end{cases} \quad (19.21)$$

Benchmarks and Toy Problems

In this chapter, we discuss some benchmark and toy problems which are used to demonstrate the utility of global optimization algorithms. Such problems have usually no direct real-world application but are well understood, widely researched, and can be used

- to measure the speed and the ability of genetic algorithm algorithms to find good solutions, as done for example by Luke and Panait [1505], Borgulya [474], and Liu et al. [1506],
- as benchmark for comparing different optimization approaches, as done by Zitzler et al. [1507] and Purshouse and Fleming [1508], for instance,
- to derive theoretical results since they are normally well understood in a mathematical sense, as done for example by Jansen and Wegener [1509],
- as basis to verify theories, as used for instance by Burke et al. [1510] and Langdon and Poli [1111],
- as playground to test new ideas, research, and developments,
- as easy-to-understand examples to discuss features and problems of optimization (as done here in Section 1.2.2 on page 9),
- for demonstration purposes, since they normally are interesting, funny, and can be visualized in a nice manner.

20.1 Real Problem Spaces

Mathematical benchmark functions are especially interesting for testing and comparing techniques based on real vectors ($\mathbb{X} = \mathbb{R}^n$) like plain Evolution Strategy (see Chapter 5 on page 191), Differential Evolution (see Section 5.5 on page 193), and Particle Swarm Optimization (see Chapter 14 on page 237). However, they only require such vectors as solution candidates, i. e., elements of the problem space $\mathbb{X}$. Hence, techniques with different search spaces $\mathbb{G}$, like genetic algorithms, can also be applied to them, given that a genotype-phenotype mapping is provided accordingly.

The optima or the Pareto frontier of benchmark functions has already been determined theoretically. When applying an optimization algorithm to the functions, we are interested in the number of solution candidates which they need to process in order to find the optima and how close we can get to them. They also give us a great opportunity to find out about the influence of parameters like population size, the choice of the selection algorithm, or the efficiency of reproduction operations.

20.1.1 Single-Objective Optimization

In this section, we list some of the most important benchmark functions for scenarios involving only a single optimization criterion. This, however, does not mean that the search

space has only a single dimension – even a single-objective optimization can take place in n -dimensional space $\mathbb{R}^n$.

Sphere

The sphere function listed by Suganthan et al. [1511] (or F_1 by De Jong [254]) and defined here in Table 20.1 is a very simple measure of efficiency of optimization methods. They have, for instance, been used by Rechenberg [105] for testing his Evolution Strategy-approach.

function	$f_{sphere}(x) = \sum_{i=1}^n x_i^2$	(20.1)
domain	$\mathbb{X} \subset \mathbb{R}^n, X_i \in [-10, 10]$	(20.2)
optimum	$\mathbf{x}^* = (0, 0, \dots, 0)^T$	(20.3)
separable	yes	
multimodal	no	

Table 20.1: The Sphere function.

TODO

20.1.2 Multi-Objective Optimization

In this section, we list some of the most important benchmark functions for scenarios involving multiple objectives (see Section 1.2.2 on page 9).

TODO

20.1.3 Dynamic Fitness Landscapes

The moving peaks benchmarks independently developed by Branke [227] and Morrison and De Jong [225] in order to illustrate the behavior of dynamic environments as discussed in Section 1.4.8 on page 53. Figure 20.1 shows an example of this benchmark for a two-dimensional real parameter setting (the third dimension is the fitness).









Fig. 20.1.1: $t = 0$ Fig. 20.1.2: $t = 1$ Fig. 20.1.3: $t = 2$ Fig. 20.1.4: $t = 3$ Fig. 20.1.5: $t = 4$ Fig. 20.1.6: $t = 6$ Fig. 20.1.7: $t = 7$ Fig. 20.1.8: $t = 13$

Fig. 20.1: An example for the moving peaks benchmark of Branke [227]

20.2 Binary Problem Spaces

20.2.1 Kauffman's NK Fitness Landscapes

The ideas of fitness landscapes¹ and epistasis² came originally from evolutionary biology and later were adopted by evolutionary computation theorists. It is thus not surprising that biologists also contributed much to the research of both. In the late 1980s, Kauffman [103] defined the *NK fitness landscape* [1512, 103, 1513], a family of objective functions with tunable epistasis, in an effort to investigate the links between epistasis and ruggedness.

The problem space and also the search space of this problem are bit strings of the length N , i. e., $\mathbb{G} = \mathbb{X} = \mathbb{B}^N$. Only one single objective function is used and referred to as fitness

¹ Fitness landscapes have been introduced in Section 1.3.2 on page 29.

² Epistasis is discussed in Section 1.4.5 on page 47.

function $F_{N,K} : \mathbb{B}^N \mapsto \mathbb{R}^+$. Each gene x_i contributes one value $f_i : \mathbb{B}^{K+1} \mapsto [0, 1] \subset \mathbb{R}^+$ to the fitness function which is defined as the average of all of these N contributions. The fitness f_i of a gene x_i is determined by its allele and the alleles at K other loci $x_{i_1}, x_{i_2}, \dots, x_{i_K}$ with $i_{1..K} \in [0, N-1] \setminus \{i\} \subset \mathbb{N}_0$, called its *neighbors*.

$$F_{N,K}(x) = \frac{1}{N} \sum_{i=0}^{N-1} f_i(x_i, x_{i_1}, x_{i_2}, \dots, x_{i_K}) \quad (20.4)$$

Whenever the value of a gene changes, all the fitness values of the genes to whose neighbor set it belongs will change too – to values uncorrelated to their previous state. While N describes the basic problem complexity, the intensity of this epistatic effect can be controlled with the parameter $K \in 0..N$: If $K = 0$, there is no epistasis at all, but for $K = N - 1$ the epistasis is maximized and the fitness contribution of each gene depends on all other genes. Two different models are defined for choosing the K neighbors: *adjacent neighbors*, where the K nearest other genes influence the fitness of a gene or *random neighbors* where K other genes are therefore randomly chosen.

The single functions f_i can be implemented by a table of length 2^{K+1} which is indexed by the (binary encoded) number represented by the gene x_i and its neighbors. These tables contain a fitness value for each possible value of a gene and its neighbors. They can be filled by sampling an uniform distribution in $[0, 1)$ (or any other random distribution).

We may also consider the f_i to be single objective functions that are combined to a fitness value $F_{N,K}$ by a weighted sum approach, as discussed in Section 1.2.3. Then, the nature of NK problems will probably lead to another well known aspect of multi-objective optimization: conflicting criteria. An improvement in one objective may very well lead to degeneration in another one.

The properties of the NK landscapes have intensely been studied in the past and the most significant results from Kauffman [49], Weinberger [1514], and Fontana et al. [1515] will be discussed here. We therefore borrow from the summaries provided by Altenberg [186] and Defoin Platel et al. [1516].

$K = 0$

For $K = 0$, the fitness function is not epistatic. Hence, all genes can be optimized separately and we have the classical additive multi-locus model.

- There is a single optimum $\mathbf{x}^*$ which is globally attractive, i. e., which can and will be found by any (reasonable) optimization process regardless of the initial configuration.
- For each individual $x \neq \mathbf{x}^*$, there exists a fitter neighbor.
- An adaptive walk³ from any point in the search space will proceed by reducing the Hamming distance to the global optimum by 1 in each step (if each mutation only affects one single gene). The number of better neighbors equals the Hamming distance to the global optimum. Hence, the estimated number of steps of such a walk is $\frac{N}{2}$.
- The fitness of direct neighbors is highly correlated since it shares $N - 1$ components.

$K = N - 1$

For $K = N - 1$, the fitness function equals a random assignment of fitness to each point of the search space.

- The probability that a genotype is a local optimum is $\frac{1}{N-1}$.
- The expected total number of local optima is thus $\frac{2^N}{N+1}$.
- The average distance between local optima is approximately $2 \ln(N - 1)$.
- The expected length of adaptive walks is approximately $\ln(N - 1)$.

³ See Section 16.4.3 on page 252 for a discussion of adaptive walks.

- The expected number of mutants to be tested in an adaptive walk before reaching a local optimum is $\sum_{i=0}^{\log_2(N-1)-1} 2^i$.
- With increasing N , the expected fitness of local optima reached by an adaptive from a random initial configuration decreases towards the mean fitness $\bar{F}_{N,K} = \frac{1}{2}$ of the search space. This is called the *complexity catastrophe* [49].

For $K = N - 1$, the work of Flyvbjerg and Lautrup [1517] is of further interest.

Intermediate K

- For small K , the best local optima share many common alleles. As K increases, this correlation diminishes. This degeneration proceeds faster for the random neighbors method than for the nearest neighbors approach.
- For larger K , the fitness of the local optima approach a normal distribution with mean m and variance s approximately

$$m = \mu + \sigma \sqrt{2 \ln(K+1)} K + 1 \quad (20.5)$$

$$s = \frac{(K+1)\sigma^2}{N(K+1+2(K+2)\ln(K+1))} \quad (20.6)$$

where μ is the expected value of the f_i and σ^2 is their variance.

- The mean distance between local optima, roughly twice the length of an adaptive walk, is approximately $\frac{N \log_2(K+1)}{2(K+1)}$.
- The autocorrelation function⁴ $\rho(k, F_{N,K})$ and the correlation length τ are:

$$\rho(k, F_{N,K}) = \left(1 - \frac{K+1}{N}\right)^k \quad (20.7)$$

$$\tau = \frac{-1}{\ln\left(1 - \frac{K+1}{N}\right)} \quad (20.8)$$

Computational Complexity

Altenberg [186] nicely summarizes the four most important theorems about the computational complexity of optimization of NK fitness landscapes. These theorems have been proven using different algorithms introduced by Weinberger [1518] and Thompson and Wright [1519].

- The NK optimization problem with adjacent neighbors is solvable in $\mathbf{O}(2^K N)$ steps and thus in $\mathcal{P}$ [1518].
- The NK optimization problem with random neighbors is $\mathcal{NP}$ -complete for $K \geq 2$ [1518, 1519].
- The NK optimization problem with random neighbors and $K = 1$ is solvable in polynomial time. [1519].

Adding Neutrality – NKp, NKq, and Technological Landscapes

As we have discussed in Section 1.4.4, natural genomes exhibit a certain degree of neutrality. Therefore, researchers have proposed extensions for the NK landscape which introduce neutrality, too [1520, 1521]. Two of them, the NKp and NKq landscapes, achieve this by altering the contributions f_i of the single genes to the total fitness. In the following, assume that there are N tables, each with 2^K entries representing these contributions.

⁴ See Definition 1.53 on page 43 for more information on autocorrelation.

The NKp landscapes devised by Barnett [145] achieves neutrality by setting a certain number of entries in each table to zero. Hence, the corresponding allele combinations do not contribute to the overall fitness of an individual. If a mutation leads to a transition from one such zero configuration to another one, it is effectively neutral. The parameter p denotes the probability that a certain allele combination does not contribute to the fitness. As proven by Reidys and Stadler [1522], the ruggedness of the NKp landscape does not vary for different values of p . Barnett [169] proved that the degree of neutrality in this landscape depends on p .

Newman and Engelhardt [1523] follow a similar approach with their NKq model. Here, the fitness contributions f_i are integers drawn from the range $[0, q)$ and the total fitness of a solution candidate is normalized by multiplying it with $1/(q-1)$. A mutation is neutral when the new allelic combination resulting from it leads to the same contribution than the old one. In NKq landscapes, the neutrality decreases with rising values of q . In [1521], you can find a thorough discussion of the NK, the NKp, and the NKq fitness landscape.

With their technological landscapes, Lobo et al. [1524] follow the same approach from the other side: they discretize the continuous total fitness function $F_{N,K}$. The parameter M of their technological landscapes corresponds to a number of bins $[0, 1/M)$, $[1/M, 2/M)$, $\dots$, into which the fitness values are sorted and put away.

20.2.2 The p-Spin Model

Motivated by the wish of researching the models for the origin of biological information by Anderson [1525], [1526] and Tsallis and Ferreira [1527], Amtrano et al. [1528] developed the p -spin model. This model is an alternative to the NK fitness landscape for tunable ruggedness [1529]. Other than the previous models, it includes a complete definition for all genetic operations which will be discussed in this section.

The p -spin model works with a fixed population size ps of individuals of an also fixed length N . There is no distinction between genotypes and phenotype, in other words, $\mathbb{G} = \mathbb{X}$. Each gene of an individual x is a binary variable which can take on the values -1 and 1 .

$$\mathbb{G} = \{-1, 1\}^N, \quad x_{i[j]} \in \{-1, 1\} \quad \forall i \in [1..ps], j \in [0..N-1] \quad (20.9)$$

On the 2^N possible genotypic configurations, a space with the topology of an N -dimensional hypercube is defined where neighboring individuals differ in exactly one element. On this genome, the Hamming distance dist_{Ham} can be defined as

$$\text{dist}_{Ham}(x_1, x_2) = \frac{1}{2} \sum_{i=0}^{N-1} (1 - x_{1[i]}x_{2[i]}) \quad (20.10)$$

Two configurations are said to be ν mutations away from each other if they have the Hamming distance ν . Mutation in this model is applied to a fraction of the $N * ps$ genes in the population. These genes are chosen randomly and their state is changed, i. e., $x[i] \rightarrow -x[i]$.

The objective function f_K (which is called *fitness function* in this context) is subject to maximization, i. e., individuals with a higher value of f_K are less likely to be removed from the population. For very subset z of exactly K genes of a genotype, one contribution $A(z)$ is added to the fitness.

$$f_K(x) = \sum_{\forall z \in \mathcal{P}([0..K]) \wedge |z|=K} A(x[z[0]], x[z[1]], \dots, x[z[K-1]]) \quad (20.11)$$

$A(z) = a_z * z[0] * z[1] * \dots * z[K-1]$ is the product of an evenly distributed random number a_z and the elements of z . For $K = 2$, f_2 can be written as $f_2(x) = \sum_{i=0}^{K-1} \sum_{j=0}^{K-1} a_{ij} x[i]x[j]$, which corresponds to the *spin-glass* [53, 54] function first mentioned by Anderson [1525] in this context. With rising values of K , this fitness landscape becomes more rugged. Its correlation length τ is approximately $N/2K$, as discussed thoroughly by Weinberger and Stadler [1529].

For selection, Amitrano et al. [1528] suggest to use the measure $P_D(x)$ defined by Rokhsar et al. [1526] as follows:

$$P_D(x) = \frac{1}{1 + e^{\beta(f_K(x) - H_0)}} \quad (20.12)$$

where the coefficient β is a sharpness parameter and H_0 is a threshold value. For $\beta \rightarrow \infty$, all individuals x with $f_K(x) < H_0$ will die and for $\beta = 0$, the death probability is always $\frac{1}{2}$. The individuals which have died are then replaced with copies of the survivors.

20.2.3 The ND Family of Fitness Landscapes

The ND family of fitness landscape has been developed by Beaudoin et al. [154] in order to provide a model problem with tunable neutrality.

The degree of neutrality ν is defined as the number (or, better, the fraction of) neutral neighbors (i.e., those with same fitness) of a solution candidate, as specified in Equation 1.43 on page 44. The populations of optimization processes residing on a neutral network (see Section 1.4.4 on page 45) tend to converge into the direction of the individual which has the highest degree of neutrality on it. Therefore, Beaudoin et al. [154] create a landscape with a predefined neutral degree distribution.

The search space is again the set of all binary strings of the length N , $\mathbb{G} = \mathbb{X} = \mathbb{B}^N$. Thus, a genotype has minimally 0 and at most N neighbors with Hamming distance 1 that have the same fitness. The array D has the length $N + 1$ and the element $D[i]$ represents the fraction of genotypes in the population that have i neutral neighbors.

Beaudoin et al. [154] provide an algorithm that divides the search space into neutral networks according to the values in D . Since this approach cannot exactly realize the distribution defined by D , the degrees of neutrality of the single individuals are subsequently refined with a Simulated Annealing algorithm. The objective (fitness) function is created in form of a complete table mapping $\mathbb{X} \mapsto \mathbb{R}$. All members of a neutral network then receive the same, random fitness.

If it is ensured that all members in a neutral network always have the same fitness, its actual value can be modified without changing the topology of the network. Tunable deceptiveness is achieved by setting the fitness values according to the *Trap* Functions [1530, 1531, 122].

Trap Functions

Trap functions $f_{b,r,\mathbf{x}^*} : \mathbb{B}^N \mapsto \mathbb{R}$ are subject to maximization based on the Hamming distance to a pre-defined global optimum $\mathbf{x}^*$. They build a second, local optimum in form of a hill with a gradient pointing away from the global optimum. This trap is parameterized with two values, b and r , where b corresponds to the width of the attractive basins and r to their relative importance.

$$f_{b,r,\mathbf{x}^*}(x) = \begin{cases} 1 - \frac{\text{dist}_{Ham}(x, \mathbf{x}^*)}{Nb} & \text{if } N * \text{dist}_{Ham}(x, \mathbf{x}^*) < b \\ \frac{r(\frac{1}{N}\text{dist}_{Ham}(x, \mathbf{x}^*) - b)}{1-b} & \text{otherwise} \end{cases} \quad (20.13)$$

20.2.4 The Royal Road

The Royal Road functions developed by Mitchell, Forrest, and Holland [781] and presented first at the Fifth International Conference on Genetic Algorithms in July 1993 are a set of special fitness landscapes for genetic algorithms [1532, 781, 1533, 1534, 161]. Their problem space $\mathbb{X}$ and search space $\mathbb{G}$ are fixed-length bit strings. The Royal Road functions are closely related to the Schema Theorem⁵ and the Building Block Hypothesis⁶ and were used to study

⁵ See Section 3.7 on page 123 for more details.

⁶ The Building Block Hypothesis is elaborated on in Section 3.7.5 on page 125

the way in which highly fit schemas are discovered. They therefore define a set of schemas $S = s_1, s_2, \dots, s_n$ and an objective function (here referred to as fitness function), subject to maximization, as

$$f(x) = \sum_{\forall s \in S} c(s) \sigma(s, x) \quad (20.14)$$

where $x \in \mathbb{X}$ is a bit string, $c(s)$ is a value assigned to the schema s and $\sigma(s, x)$ is defined as

$$\sigma(s, x) = \begin{cases} 1 & \text{if } x \text{ is an instance of } s \\ 0 & \text{otherwise} \end{cases} \quad (20.15)$$

In the original version, $c(s)$ is the order of the schema s , i. e., $c(s) \equiv \text{order}(s)$, and S is specified as follows (where $*$ stands for the *don't care* symbol as usual).

1	S_1	=	11111111*****	$c(S_1) = 8$
2	S_2	=	*****1111111*****	$c(S_2) = 8$
3	S_3	=	*****1111111*****	$c(S_3) = 8$
4	S_4	=	*****1111111*****	$c(S_4) = 8$
5	S_5	=	*****1111111*****	$c(S_5) = 8$
6	S_6	=	*****1111111*****	$c(S_6) = 8$
7	S_7	=	*****1111111*****	$c(S_7) = 8$
8	S_8	=	*****1111111*****	$c(S_8) = 8$
9	S_9	=	111111111111*****	$c(S_9) = 16$
10	S_{10}	=	*****111111111111*****	$c(S_{10}) = 16$
11	S_{11}	=	*****111111111111*****	$c(S_{11}) = 16$
12	S_{12}	=	*****111111111111*****	$c(S_{12}) = 16$
13	S_{13}	=	1111111111111111*****	$c(S_{13}) = 32$
14	S_{14}	=	*****1111111111111111*****	$c(S_{14}) = 32$
15	S_{15}	=	11111111111111111111*****	$c(S_{15}) = 64$

Listing 20.1: An example Royal Road function.

By the way, from this example, we can easily see that a fraction of all mutation and crossover operations applied to most of the solution candidates will fall into the don't care areas. Such modifications will not yield any fitness change and therefore are neutral.

The Royal Road functions provide certain, predefined stepping stones (i. e., building blocks) which (theoretically) can be combined by the genetic algorithm to successively create schemas of higher fitness and order. Mitchell et al. [781] performed several tests with their Royal Road functions. These tests revealed or confirmed that

- Crossover is a useful reproduction operation in this scenario. Genetic algorithms which apply this operation clearly outperform hill climbing approaches solely based on mutation.
- In the spirit of the Building Block Hypothesis, one would expect that the intermediate steps (for instance order 32 and 16) of the Royal Road functions would help the genetic algorithm to reach the optimum. The experiments of Mitchell et al. [781] showed the exact opposite: leaving them away speeds up the evolution significantly. The reason is the fitness difference between the intermediate steps and the low-order schemas is high enough that the first instance of them will lead the GA to converge to it and wipe out the low-order schemas. The other parts of this intermediate solution play no role and may allow many zeros to *hitchhike* along.

Especially this last point gives us another insight on how we should construct genomes: the fitness of combinations of good low-order schemas should not be too high so other good low-order schemas do not extinct when they emerge. Otherwise, the phenomenon of *domino convergence* researched by Rudnick [75] and outlined in Section 1.4.1 and Section 20.2.5 may occur.

Variable-Length Representation

The original Royal Road problems can be defined for binary string genomes of any given length n , as long as n is fixed. A Royal Road benchmark for variable-length genomes has been defined by Defoin Platel et al. [1535].

The problem space $\mathbb{X}_\Sigma$ of the VLR (variable-length representation) Royal Road problem is based on an alphabet Σ with $N = |\Sigma|$ letters. The fitness of an individual $x \in \mathbb{X}_\Sigma$ is determined by whether or not consecutive *building blocks* of the length b of the letters $l \in \Sigma$ are present. This presence can be defined as

$$B_b(x, l) = \begin{cases} 1 & \text{if } \exists i : (0 \leq i < (\text{len}(x) - b)) \wedge (x[i+j] = l \forall j : 0 \leq j < (b - 1)) \\ 0 & \text{otherwise} \end{cases} \quad (20.16)$$

- Where $b \geq 1$ is the length of the building blocks,
- Σ is the alphabet with $N = |\Sigma|$ letters,
- l is a letter in Σ ,
- $x \in \mathbb{X}_\Sigma$ is a solution candidate, and
- $x[k]$ is the k^{th} locus of x .

$B_b(x, l)$ is 1 if a building block, an uninterrupted sequence of the letter l , of at least length b , is present in x . Of course, if $\text{len}(x) < b$ this cannot be the case and $B_b(x, l)$ will be zero.

We can now define the functional objective function $f_{\Sigma b} : \mathbb{X}_\Sigma \mapsto [0, 1]$ which is subject to maximization as

$$f_{\Sigma b}(x) = \frac{1}{N} \sum_{i=1}^N B_b(x, \Sigma[i]) \quad (20.17)$$

An optimal individual $\mathbf{x}^*$ solving the VLR Royal Road problem is thus a string that includes building blocks of length b for all letters $l \in \Sigma$. Notice that the position of these blocks plays no role. The set $\mathbf{X}_b^*$ of all such optima with $f_{\Sigma b}(\mathbf{x}^*) = 1$ is then

$$\mathbf{X}_b^* \equiv \{\mathbf{x}^* \in \mathbb{X}_\Sigma : B_b(\mathbf{x}^*, l) = 1 \forall l \in \Sigma\} \quad (20.18)$$

Such an optimum $\mathbf{x}^*$ for $b = 3$ and $\Sigma = \{A, T, G, C\}$ is

$$\mathbf{x}^* = \mathbf{AAAGTGGGTAATTTTCCCTCCC} \quad (20.19)$$

The relevant building blocks of $\mathbf{x}^*$ are written in bold face. As it can easily be seen, their location plays no role, only their presence is important. Furthermore, multiple occurrences of building blocks (like the second *CCC*) do not contribute to the fitness. The fitness landscape has been designed in a way ensuring that fitness degeneration by crossover can only occur if the crossover points are located inside building blocks and not by block translocation or concatenation. In other words, there is no inter-block epistasis.

Epistatic Road

Defoin Platel et al. [1516] combined their previous work on the VLR Royal Road with Kauffman's NK landscapes and introduced the Epistatic Road. The original NK landscape works on binary representation of the fixed length N . To each locus i in the representation, one fitness function f_i is assigned denoting its contribution to the overall fitness. f_i however is not exclusively computed using the allele at the i^{th} locus but also depends on the alleles of K other loci, its neighbors.

The VLR Royal Road uses a genome based on the alphabet Σ with $N = \text{len}(\Sigma)$ letters. It defines the function $B_b(x, l)$ which returns 1 if a building block of length b containing only the character l is present in x and 0 otherwise. Because of the fixed size of the alphabet Σ , there exist exactly N such functions. Hence, the variable-length representation can be translated to a fixed-length, binary one by simply concatenating them:

$$B_b(x, \Sigma_{[0]}) B_b(x, \Sigma_{[1]}) \dots B_b(x, \Sigma_{[N-1]}) \quad (20.20)$$

Now we can define a NK landscape for the Epistatic Road by substituting the $B_b(x, l)$ into Equation 20.4 on page 282:

$$F_{N,K,b}(x) = \frac{1}{N} \sum_{i=0}^{N-1} f_i(B_b(x, \Sigma_{[i]}), B_b(x, \Sigma_{[i_1]}), \dots, B_b(x, \Sigma_{[i_K]})) \quad (20.21)$$

The only thing left is to ensure that the end of the road, i. e., the presence of all N building blocks, also is the optimum of $F_{N,K,b}$. This is done by exhaustively searching the space $\mathbb{B}^N$ and defining the f_i in a way that $B_b(x, l) = 1 \forall l \in \Sigma \Rightarrow F_{N,K,b}(x) = 1$.

Royal Trees

An analogue of the Royal Road for Genetic Programming has been specified by Punch et al. [1536]. This *Royal Tree* problem specifies a series of functions $A, B, C, \dots$ with increasing arity, i. e., A has one argument, B has two arguments, C has three, and so on. Additionally, a set of terminal nodes x, y, z is defined.


Fig. 20.2.1: Perfect A -levelFig. 20.2.2: Perfect B -levelFig. 20.2.3: Perfect C -level

Fig. 20.2: The perfect Royal Trees.

For the first free levels, the perfect trees are shown Figure 20.2. An optimal A -level tree consists of an A node with an x leaf attached to it. The perfect level- B tree has a B as root with two perfect level- A trees as children. A node labeled with C having three children which all are optimal B -level trees is the optimum at C -level, and so on.

The objective function, subject to maximization, is computed recursively. The raw fitness of a node is the weighted sum of the fitness of its children. If the child is a perfect tree at the appropriate level, a perfect C tree beneath a D -node, for instance, its fitness is multiplied with the constant *FullBonus*, which normally has the value 2. If the child is not a perfect tree, but has the correct root, the weight is *PartialBonus* (usually 1). If it is otherwise incorrect, its fitness is multiplied with *Penalty*, which is $\frac{1}{3}$ per default. If the whole tree is a perfect tree, its raw fitness is finally multiplied with *CompleteBonus* which normally is also 2. The value of a x leaf is 1.

From Punch et al. [1536], we can furthermore borrow three examples for this fitness assignment and outline them in Figure 20.3. A tree which represents a perfect A level has the score of $CompleteBonus * FullBonus * 1 = 2 * 2 * 1 = 4$. A complete and perfect tree at level B receives $CompleteBonus(FullBonus * 4 + FullBonus * 4) = 2 * (2 * 4 + 2 * 4) = 32$. At level C , this makes $CompleteBonus(FullBonus * 32 + FullBonus * 32 + FullBonus * 32) = 2(2 * 32 + 2 * 32 + 2 * 32) = 384$.

Other Derived Problems

Storch and Wegener [1537; 1538; 1539] used their Real Royal Road for showing that there exist problems where crossover helps improving the performance of evolutionary algorithms.




Fig. 20.3.1: $2(2 * 32 + 2 * 32 + 2 * 32) = 384$




Fig. 20.3.2: $2(2 * 32 + 2 * 32 + \frac{2}{3} * 1) = 128\frac{2}{3}$




Fig. 20.3.3: $2(2 * 32 + \frac{1}{3} * 1 + \frac{1}{3} * 1) = 64\frac{2}{3}$

Fig. 20.3: Example fitness evaluation of Royal Trees

20.2.5 OneMax and BinInt

The OneMax and BinInt are two very simple model problems for measuring the convergence of genetic algorithms.

The OneMax Problem

The task in the OneMax (or BitCount) problem is to find a binary string of length n consisting of all ones. The search and problem space are both the fixed-length bit strings $\mathbb{G} = \mathbb{X} = \mathbb{B}^n$. Each gene (bit) has two alleles 0 and 1 which also contribute exactly this value to the total fitness, i. e.,

$$f(x) = \sum_{i=0}^{n-1} x[i], \quad \forall x \in \mathbb{X} \quad (20.22)$$

For the OneMax problem, an extensive body of research has been provided by Mühlenbein and Schlierkamp-Voosen [1540], Thierens and Goldberg [1541], Miller and Goldberg [220], Bäck [1542], Blickle and Thiele [648], and Wilson and Kaur [1543].

The BinInt Problem

The BinInt problem devised by Rudnick [75] also uses the bit strings of the length n as search and problem space ($\mathbb{G} = \mathbb{X} = \mathbb{B}^n$). It is something like a perverted version of the OneMax problem, with the objective function defined as

$$f(x) = \sum_{i=0}^{n-1} 2^{n-i-1} x[i], \quad x[i] \in \{0, 1\} \quad \forall i \in [0..n-1] \quad (20.23)$$

Since the bit at index i has a higher contribution to the fitness than all other bit at higher indices together, the comparison between two solution candidates x_1 and x_2 is won by the lexicographically bigger one. Thierens et al. [76] give the example $x_1 = (1, 1, \underline{1}, 1, 0, 0, 0, 0)$ and $x_2 = (1, 1, \underline{0}, 0, 1, 1, 0, 0)$, where the first deviating bit (underlined, at index 2) fully determines the outcome of the comparison of the two.

We can expect that the bits with high contribution (high salience) will converge quickly whereas the other genes with lower salience are only pressured by selection when all others have already been fully converged. Rudnick [75] called this sequential convergence phenomenon *domino convergence* due to its resemblance with a row of falling domino stones [76] (see Section 1.4.1). Generally, he showed that first, the highly salient genes converge (i. e., take on the correct values in the majority of the population). Then, step by step, the building blocks of lower significance can converge, too. Another result of Rudnick's work

is that mutation may stall the convergence because it can disturb the highly significant genes, which then counters the effect of the selection pressure on the less salient ones. Then, it becomes very less likely that the majority of the population will have the best alleles in these genes. This somehow dovetails with the idea of error thresholds from theoretical biology [130, 131] which we have mentioned in Section 1.4.2. It also explains some of the experimental results obtained with the Royal Road problem from Section 20.2.4. The BinInt problem was used in the studies of Sastry and Goldberg [1544; 1545].

20.2.6 Tunable Model for Problematic Phenomena

What is a good model problem? Which model fits best to our purposes? These questions should be asked whenever we apply a benchmark, whenever we want to use something for testing the ability of a global optimization approach. The mathematical functions introduced in Section 20.1, for instance, are good for testing special mathematical reproduction operations like used in Evolution Strategies and for testing the capability of an evolutionary algorithm for estimating the Pareto frontier in multi-objective optimization. Kauffman's NK fitness landscape (discussed in Section 20.2.1) was intended to be a tool for exploring the relation of ruggedness and epistasis in fitness landscapes but can prove very useful for finding out how capable an global optimization algorithm is to deal with problems exhibiting these phenomena. In Section 20.2.4, we outlined the Royal Road functions, which were used to investigate the ability of genetic algorithms to combine different useful formae and to test the Building Block Hypothesis. The Artificial Ant (Section 20.3.1) and the GCD problem from Section 20.3.2 are tests for the ability of Genetic Programming of learning algorithms.

All these benchmarks and toy problems focus on specific aspects of global optimization and will exhibit different degrees of the problematic properties of optimization problems to which we had devoted Section 1.4:

- premature convergence and multimodality (Section 1.4.1),
- ruggedness (Section 1.4.2),
- deceptiveness (Section 1.4.3),
- neutrality and redundancy (Section 1.4.4),
- overfitting and oversimplification (Section 1.4.6), and
- dynamically changing objective functions (Section 1.4.8).

With the exception of the NK fitness landscape, it remains unclear to which degrees these phenomena occur in the test problem. How much intrinsic epistasis does the Artificial Ant or the GCD problem emit? What is the quantity of neutrality of Royal Road for variable-length representations emit? Are the mathematical test functions rugged and, if so, to which degree? All the problems are useful test instances for global optimization. They have not been designed to give us answers to questions like: Which fitness assignment process can be useful when an optimization problem exhibits weak causality and thus has a rugged fitness landscape? How does a certain selection algorithm influence the ability of a genetic algorithm to deal with neutrality? Only Kauffman's NK landscape provides such answers, but only for epistasis. By fine-tuning its N and K parameters, we can generate problems with different degrees of epistasis. Applying a genetic algorithm to these problems then allows us to draw conclusions on its expected performance when being fed with high or low epistatic real-world problems.

In this section, a new model problem is defined that exhibits ruggedness, epistasis, neutrality, multi-objectivity, overfitting, and oversimplification features in a controllable manner Weise et al. [1546], Niemczyk [1547]. Each of them is introduced as a distinct filter component which can separately be activated, deactivated, and fine-tuned. This model provides a perfect test bed for optimization algorithms and their configuration settings. Based on a rough estimation of the structure of the fitness landscape of a given problem, tests can be run very fast using the model as a benchmark for the settings of an optimization algorithm.

Thus, we could, for instance, determine a priori whether increasing the population size of an evolutionary algorithm over an approximated limit is likely to provide a gain.

With it, we also can evaluate the behavior of an optimization method in the presence of various problematic aspects, like epistasis or neutrality. This way, strengths and weaknesses of different evolutionary approaches could be explored in a systematic manner. Additionally, it is also well suited for theoretical analysis because of its simplicity. The layers of the model, sketched using an example in Figure 20.4, are specified in the following.




Fig. 20.4: An example for the fitness landscape model.

Model Definition

The basic optimization task in this model is to find a binary string $\mathbf{x}^*$ of a predefined length $n = \text{len}(\mathbf{x}^*)$ consisting of alternating zeros and ones in the space of all possible binary strings $\mathbb{X} = \mathbb{B}^*$. The tuning parameter for the problem size is $n \in \mathbb{N}$.

$$\mathbf{x}^* = 0101010101010 \dots 01 \quad (20.24)$$

Overfitting and Oversimplification

Searching this optimal string could be done by comparing each genotype g with $\mathbf{x}^*$. Therefore we would use the Hamming distance⁷ [627] $\text{dist}_{Ham}(a, b)$, which defines the difference between two binary strings a and b of equal length as the number of bits in which they differ (see Equation 28.10).

Instead of doing this directly, we test the solution candidate against tc training samples $T_1, T_2, \dots, T_{tc}$. These samples are modified versions of the perfect string $\mathbf{x}^*$.

As outlined in Section 1.4.6 on page 48, we can distinguish between overfitting and oversimplification. The latter is often caused by incompleteness of the training cases and the former can originate from noise in the training data. Both forms can be expressed in terms of our model by the objective function $f_{\varepsilon, o, tc}$ (based on a slightly extended version of the Hamming distance dist_{Ham}^*) which is subject to minimization.

$$\text{dist}_{Ham}^*(a, b) = |\{\forall i : (a[i] \neq b[i]) \wedge (b[i] \neq *) \wedge (0 \leq i < |a|)\}| \quad (20.25)$$

$$f_{\varepsilon, o, tc}(x) = \sum_{i=1}^{tc} \text{dist}_{Ham}^*(x, T_i), \quad f_{\varepsilon, o, tc}(x) \in [0, \hat{f}] \quad \forall x \in \mathbb{X} \quad (20.26)$$

In the case of oversimplification, the perfect solution $\mathbf{x}^*$ will always reach a perfect score in all training cases. There may be incorrect solutions reaching this value in some cases too, because some of the facets of the problem are hidden. We take this into consideration by placing *o don't care* symbols ($*$) uniformly distributed into the training cases. The values of the solution candidates at their loci have no influence on the fitness.

When overfitting is enabled, the perfect solution will not reach the optimal score in any training case because of the noise present. Incorrect solutions may score better in some cases and even outperform the real solution if the noise level is high. Noise is introduced in the training cases by toggling ε of the remaining $n - o$ bits, again following a uniform distribution. An optimization algorithm can find a correct solution only if there are more training samples with correctly defined values for each locus than with wrong or don't care values.

The optimal objective value is zero and the maximum $\hat{f}$ of $f_{\varepsilon, o, tc}$ is limited by the upper boundary $\hat{f} \leq (n - o)tc$. Its exact value depends on the training cases. For each bit index i , we have to take into account whether a zero or a one in the phenotype would create larger errors:

$$\text{count}(i, val) = |\{j \in 1..n : T_j[i] = val\}| \quad (20.27)$$

$$e(i) = \begin{cases} \text{count}(i, 1) & \text{if } \text{count}(i, 1) \geq \text{count}(i, 0) \\ \text{count}(i, 0) & \text{otherwise} \end{cases} \quad (20.28)$$

$$\hat{f} = \sum_{i=1}^{tc} e(i) \quad (20.29)$$

⁷ Definition 28.6 on page 469 includes the specification of the Hamming distance.

Neutrality

We can create a well-defined amount of neutrality during the genotype-phenotype mapping by applying a transformation u_μ that shortens the solution candidates by a factor μ . The i^{th} bit in $u_\mu(g)$ is defined as 0 if and only if the majority of the μ bits starting at locus $i * \mu$ in g is also 0, and as 1 otherwise. The default value 1 set in draw situations has (in average) no effect on the fitness since the target solution $\mathbf{x}^*$ is defined as a sequence of alternating zeros and ones. If the length of a genotype g is not a multiple of μ , the remaining $\text{len}(g) \bmod \mu$ bits are ignored. The tunable parameter for the neutrality in our model is μ . If μ is set to 1, no additional neutrality is modeled.

Epistasis

Epistasis in general means that a slight change in one gene of a genotype influences some other genes. We can introduce epistasis in our model as part of the genotype mapping and apply it after the neutrality transformation. We therefore define a bijective function e_η that translates a binary string z of length η to a binary string $e_\eta(z)$ of the same length. Assume we have two binary strings z_1 and z_2 which differ only in one single locus, i.e., their Hamming distance is one. $e_\eta(i)$ introduces epistasis by exhibiting the following property:

$$\text{dist}_{Ham}(z_1, z_2) = 1 \Rightarrow \text{dist}_{Ham}(e_\eta(z_1), e_\eta(z_2)) \geq \eta - 1 \quad \forall z_1, z_2 \in \mathbb{B}^\eta \quad (20.30)$$

The meaning of Equation 20.30 is that a change of one bit in a genotype g leads to the change of at least $\eta - 1$ bits in the corresponding mapping $e_\eta(x)$. This, as well as the demand for bijectivity, is provided if we define e_η as follows:

$$e_\eta(z)[i] = \begin{cases} \bigotimes_{\substack{\forall j: 0 \leq j < \eta \wedge \\ j \neq (i-1) \bmod \eta}} z[j] & \text{if } 0 \leq z < 2^{\eta-1}, \text{ i. e., } z[\eta-1] = 0 \\ \overline{e_\eta(z - 2^{\eta-1})[i]} & \text{otherwise} \end{cases} \quad (20.31)$$

In other words, for all strings $z \in \mathbb{B}^\eta$ which have the most significant bit (MSB) not set, the e_η transformation is performed bitwise. The i^{th} bit in $e_\eta(z)$ equals the exclusive or combination of all but one bit in z . Hence, each bit in z influences the value of $\eta - 1$ bits in $e_\eta(z)$. For all z with 1 in the MSB, $e_\eta(z)$ is simply set to the negated e_η transformation of z with the MSB cleared (the value of the MSB is $2^{\eta-1}$). This division in e is needed in order to ensure its bijectiveness. This and the compliance with Equation 20.30 can be shown with a rather lengthy proof omitted here.

In order to introduce this model of epistasis in genotypes of arbitrary length, we divide them into blocks of the length η and transform each of them separately with e_η . If the length of a given genotype g is no multiple of η , the remaining $\text{len}(g) \bmod \eta$ bits at the end will be transformed with the function $e_{\text{len}(g) \bmod \eta}$ instead of e_η , as outlined in Figure 20.4. It may also be an interesting fact that the e_η transformations are a special case of the NK landscape discussed in Section 20.2.1 with $N = \eta$ and $K \approx \eta - 2$.

z	$e_4(z)$	z	$e_4(z)$	z	$e_4(z)$
0000	→ 0000	1111	→ 1110	0011	→ 0110
0001	→ 1101	0111	→ 0001	0101	→ 1010
0010	→ 1011	1011	→ 1001	0110	→ 1100
0100	→ 0111	1101	→ 0101	1001	→ 0010
1000	→ 1111	1110	→ 0011	1010	→ 0100
				1100	→ 1000

Fig. 20.5: An example for the epistasis mapping $z \rightarrow e_4(z)$.

The tunable parameter η for the epistasis ranges from 2 to $n * m$, the product of the basic problem length n and the number of objectives m (see next section). If it is set to a value smaller than 3, no additional epistasis is introduced. Figure 20.5 outlines the mapping for $\eta = 4$.

Multi-Objectivity

A multi-objective problem with m criteria can easily be created by interleaving m instances of the benchmark problem with each other and introducing separate objective functions for each of them. Instead of just dividing the genotype g in m blocks, each standing for one objective, we scatter the objectives as illustrated in Figure 20.4. The bits for the first objective function comprise $x_1 = (g[0], g[m], g[2m], \dots)$, those used by the second objective function are $x_2 = (g[1], g[m+1], g[2m+1], \dots)$. Notice that no bit in g is used by more than one objective function. Superfluous bits (beyond index $nm - 1$) are ignored. If g is too short, the missing bits in the phenotypes are replaced with the complement from $\mathbf{x}^*$, i. e., if one objective misses the last bit (index $n - 1$), it is padded by $\overline{\mathbf{x}^*_{[n-1]}}$ which will worsen the objective by 1 on average.

Because of the interleaving, the objectives will begin to conflict if epistasis ($\eta > 2$) is applied, similar to NK landscapes. Changing one bit in the genotype will change the outcome of at most $\min\{\eta, m\}$ objectives. Some of them may improve while others may worsen.

A non-functional objective function minimizing the length of the genotypes is added if variable-length genomes are used during the evolution. If fixed-length genomes are used, they can be designed in a way that the blocks for the single objectives have always the right length.

Ruggedness

In an optimization problem, there can be at least two (possibly interacting) sources of ruggedness of the fitness landscape. The first one, epistasis, has already been modeled and discussed. The other source concerns the objective functions themselves, the nature of the problem. We will introduce this type of ruggedness *a posteriori* by artificially lowering the causality of the problem space. We therefore shuffle the objective values with a permutation $r : [0, \hat{f}] \mapsto [0, \hat{f}]$, where $\hat{f}$ the abbreviation for the maximum possible objective value, as defined in Equation 20.29.

Before we do that, let us shortly outline what makes a function *rugged*. Ruggedness is obviously the opposite of smoothness and causality. In a smooth objective function, the objective values of the solution candidates neighboring in problem space are also neighboring. In our original problem with $o = 0$, $\varepsilon = 0$, and $tc = 1$ for instance, two individuals differing in one bit will also differ by one in their objective values. We can write down the list of objective values the solution candidates will take on when they are stepwise improved from the worst to the best possible configuration as $(\hat{f}, \hat{f} - 1, \dots, 2, 1, 0)$. If we exchange two of the values in this list, we will create some artificial ruggedness. A measure for the ruggedness of such a permutation r is $\Delta(r)$:

$$\Delta(r) = \sum_{i=0}^{\hat{f}-1} |r[i] - r[i+1]| \quad (20.32)$$

The original sequence of objective values has the minimum value $\check{\Delta} = \hat{f}$ and the maximum possible value is $\hat{\Delta} = \frac{\hat{f}(\hat{f}+1)}{2}$. There exists at least one permutation for each Δ value in $\check{\Delta}.. \hat{\Delta}$. We can hence define the permutation r_γ which is applied after the objective values are computed and which has the following features:

1. It is bijective (since it is a permutation).
2. It must preserve the optimal value, i. e., $r_\gamma[0] = 0$.

$$3. \Delta(r_\gamma) = \tilde{\Delta} + \gamma.$$

With $\gamma \in [0, \hat{\Delta} - \tilde{\Delta}]$, we can fine-tune the ruggedness. For $\gamma = 0$, no ruggedness is introduced. For a given $\hat{f}$, we can compute the permutations r_γ with the procedure $\text{buildRPermutation}(\gamma, \hat{f})$ defined in Algorithm 20.1.

Algorithm 20.1: $r_\gamma \leftarrow \text{buildRPermutation}(\gamma, \hat{f})$

Input: γ : the γ value
Input: $\hat{f}$: the maximum objective value
Data: i, j, d, tmp : temporary variables
Data: $k, start, r$: parameters of the subalgorithm
Output: r_γ : the permutation r_γ

```


1 begin
2   Subalgorithm  $r \leftarrow \text{permute}(k, r, start)$ 
3   begin
4     if  $k > 0$  then
5       if  $k \leq (\hat{f} - 1)$  then
6          $r \leftarrow \text{permute}(k - 1, r, start)$ 
7          $tmp \leftarrow r[\hat{f}]$ 
8          $r[\hat{f}] \leftarrow r[\hat{f} - k]$ 
9          $r[\hat{f} - k] \leftarrow tmp$ 
10      else
11         $i \leftarrow \lfloor \frac{start + 1}{2} \rfloor$ 
12        if  $(start \bmod 2) = 0$  then
13           $i \leftarrow \hat{f} + 1 - i$ 
14           $d \leftarrow -1$ 
15        else
16           $d \leftarrow 1$ 
17        for  $j \leftarrow start$  up to  $\hat{f}$  do
18           $r[j] \leftarrow i$ 
19           $i \leftarrow i + d$ 
20         $r \leftarrow \text{permute}(k - \hat{f} + start, r, start + 1)$ 
21      end
22       $r \leftarrow (0, 1, 2, \dots, \hat{f} - 1, \hat{f})$ 
23      return  $\text{permute}(\gamma, r, 1)$ 
24 end
```

Figure 20.6 outlines all ruggedness permutations r_γ for an objective function which can range from 0 to $\hat{f} = 5$. As can be seen, the permutations scramble the objective function more and more with rising γ and reduce its gradient information.

Experimental Validation

In this section, we will use a selection of the experimental results obtained with our model in order to validate the correctness of the approach.⁸ Table 20.2 states the configuration of the evolutionary algorithm used for our experiments. For each of the experiment-specific settings discussed later, at least 50 runs have been performed.

⁸ More experimental results and more elaborate discussions can be found in the bachelor's thesis of Niemczyk [1547].

Fig. 20.6: An example for r_γ with $\gamma = 0..10$ and $\hat{f} = 5$.

Parameter	Short	Description
Problem	$\mathbb{X}$	The variable-length bit strings consisting of between 1 and 8000 bits. (see Section 3.5)
Space		
Objective	F	$F = \{f_{\varepsilon,o,tc}, f_{nf}\}$, where f_{nf} is the non-functional length criterion
Functions		$f_{nf}(x) = \text{len}(x)$ (see Equation 20.26)
Search Space	$\mathbb{G}$	$\mathbb{G} = \mathbb{X}$
Search	Op	80% single-point crossover, 20% single-bit mutation
Operations		
GPM	gpm	(see Section 20.2.6)
Optimization	alg	plain genetic algorithm (see Chapter 3)
Algorithm		
Comparison	cm	Pareto comparison (see Section 1.2.4)
Operator		
Population	ps	$ps = 1000$
Size		
Steady-State	ss	The algorithms were generational (not steady-state) ($ss = 0$). (see Section 2.1.2)
Fitness	fa	For fitness assignment in the evolutionary algorithm, Pareto ranking was used. (see Section 2.3.3)
Assignment		
Algorithm		
Selection	sel	A tournament selection with tournament size $k = 5$ was applied. (see Section 2.4.4)
Algorithm		
Convergence	cp	No additional means for convergence prevention were used, i. e., $cp = 0$. (see Section 2.4.8)
Prevention		
Generation	mt	The maximum number of generations that each run is allowed to perform. (see Definition 1.43)
Limit		$mt = 1001$

Table 20.2: The settings of the experiments with the benchmark model.

Basic Complexity

In the experiments, we distinguish between *success* and *perfection*. Success means finding individuals x of optimal *functional* fitness, i.e., $f_{\varepsilon,o,tc}(x) = 0$. Multiple such *successful* strings may exist, since superfluous bits at the end of genotypes do not influence their functional objective. We will refer to the number of generations needed to find a successful individual as *success generations*. The perfect string $\mathbf{x}^*$ has no useless bits, it is the shortest possible

solution with $f_{\varepsilon,o,tc} = 0$ and, hence, also optimal in the non-functional length criterion. In our experiments, we measure:

Measure	Short	Description
Success Fraction	s/r	The fraction of experimental runs that turned out successful. (see Section 19.3.1)
Minimum Success Generation	$\widetilde{st}$	The number of generations needed by the fastest (successful) experimental run to find a successful individual. (see Section 19.3.1)
Mean Success Generation	$\overline{st}$	The average number of generations needed by the (successful) experimental runs to find a successful individual. (see Section 19.3.1)
Maximum Success Generation	$\widehat{st}$	The number of Generations generations needed by the slowest (successful) experimental run to find a successful individual. (see Section 19.3.1)
Mean Perfection Generation	$\overline{pt}$	The average number of generations needed by the (perfect) experimental runs to find a perfect individual. equ:experimentPerfAvgGen

Table 20.3: First-level evaluation results of the experiments with the model benchmark.




Fig. 20.7: The basic problem hardness.

In Figure 20.7, we have computed the minimum, average, and maximum number of the success generations ($\widetilde{st}$, $\overline{st}$, and $\widehat{st}$) for values of n ranging from 8 to 800. As illustrated, the problem hardness increases steadily with rising string length n . Trimming down the solution strings to the perfect length becomes more and more complicated with growing n . This is likely because the fraction at the end of the strings where the trimming is to be performed will shrink in comparison with its overall length.




Fig. 20.8: Experimental results for the ruggedness.

Ruggedness

In Figure 20.8, we plotted the average success generations $\overline{st}$ with $n = 80$ and different ruggedness settings γ . Interestingly, the gray original curve behaves very strangely. It is divided into alternating solvable and unsolvable⁹ problems. The unsolvable ranges of γ correspond to gaps in the curve. With rising γ , the solvable problems require more and more generations until they are solved. After a certain (earlier) γ threshold value, the unsolvable sections become solvable. From there on, they become simpler with rising γ . At some point, the two parts of the curve meet.

The reason for this behavior is rooted in the way that we construct the ruggedness mapping r and illustrates the close relation between ruggedness and deceptiveness. Algorithm 20.1 is a greedy algorithm which alternates between creating groups of mappings that are mainly rugged and such that are mainly deceptive. In Figure 20.6 for instance, from $\gamma = 5$ to $\gamma = 7$, the permutations exhibit a high degree of deceptiveness whilst just being rugged before and after that range. Thus, it seems to be a good idea to rearrange these sections of the ruggedness mapping. The identity mapping should still come first, followed by the purely rugged mappings ordered by their Δ -values. Then, the permutations should gradually change from rugged to deceptive and the last mapping should be the most deceptive one ($\gamma = 10$ in Figure 20.6). The black curve in Figure 20.8 depicts the results of rearranging the γ -values with Algorithm 20.2. This algorithm maps deceptive gaps to higher γ -values and, by doing so, makes the resulting curve continuous.¹⁰

Fig. 20.9.1 sketches the average success generations for the rearranged ruggedness problem for multiple values of n and γ' . Depending on the basic problem size n , the problem hardness increases steeply with rising values of γ' .

In Algorithm 20.1 and Algorithm 20.2, we use the maximum value of the functional objective function (abbreviated with $\hat{f}$) in order to build and to rearrange the ruggedness

⁹ We call a problem unsolvable if it has not been solved within 1000 generations.

¹⁰ This is a deviation from our original idea, but this idea did not consider deceptiveness.

Algorithm 20.2: $\gamma \leftarrow \text{translate}(\gamma', \hat{f})$

Input: γ' : the raw γ value
Input: $\hat{f}$: the maximum value of $f_{\varepsilon, o, tc}$
Data: i, j, k, l : some temporary variables
Output: γ : the translated γ value

```

1 begin
2    $l \leftarrow \frac{\hat{f}(\hat{f}-1)}{2}$ 
3    $i \leftarrow \left\lfloor \frac{\hat{f}}{2} \right\rfloor * \left\lfloor \frac{\hat{f}+1}{2} \right\rfloor$ 
4   if  $\gamma \leq \hat{f}i$  then
5      $j \leftarrow \left\lfloor \frac{\hat{f}+2}{2} - \sqrt{\frac{\hat{f}^2}{4} + 1 - \gamma} \right\rfloor$ 
6      $k \leftarrow \gamma - j(\hat{f} + 2) + j^2 + \hat{f}$ 
7     return  $k + 2(j(\hat{f} + 2) - j^2 - \hat{f}) - j$ 
8   else
9      $j \leftarrow \left\lfloor \frac{(\hat{f} \bmod 2) + 1}{2} + \sqrt{\frac{1 - (\hat{f} \bmod 2)}{4} + \gamma - 1 - i} \right\rfloor$ 
10     $k \leftarrow \gamma - (j - (\hat{f} \bmod 2))(j - 1) - 1 - i$ 
11    return  $l - k - 2j^2 + j - (\hat{f} \bmod 2)(-2j + 1)$ 
12 end
  
```

permutations r . Since this value depends on the basic problem length n , the number of different permutations and thus, the range of the γ' values will too. The length of the lines in direction of the γ' axis in Fig. 20.9.1 thus increases with n . We introduce two additional scaling functions for ruggedness and deceptiveness with a parameter g spanning from zero to ten, regardless of n . Only one of these functions can be used at a time, depending on whether experiments should be run for rugged (Equation 20.33) or for deceptive (Equation 20.34) problems. For scaling, we use the highest γ' value which maps to rugged mappings $\gamma'_r = \left\lfloor 0.5\hat{f} \right\rfloor * \left\lceil 0.5\hat{f} \right\rceil$, and the minimum and maximum ruggedness values according to Equation 20.32.

$$\text{rugged: } \gamma' = \text{round}(0.1g * \gamma'_r) \quad (20.33)$$

$$\text{deceptive: } \gamma' = \begin{cases} 0 & \text{if } g \leq 0 \\ \gamma'_r + \text{round}\left(0.1g * (\hat{\Delta} - \check{\Delta} - \gamma'_r)\right) & \text{otherwise} \end{cases} \quad (20.34)$$

When using this scaling mechanism, the curves resulting from experiments with different n -values can be compared more easily: Fig. 20.9.2 based on the scale from Equation 20.33, for instance, shows much clearer how the problem difficulty rises with increasing ruggedness than Fig. 20.9.1. We also can spot some irregularities which always occur at about the same degree of ruggedness, near $g \approx 9.5$, and that we will investigate in future.

The experiments with the deceptiveness scale Equation 20.34 show the tremendous effect of deceptiveness in the fitness landscape. Not only does the problem hardness rise steeply with g (Fig. 20.9.3), after certain threshold, the evolutionary algorithm becomes unable to solve the model problem at all (in 1000 generations), and the fraction of failed experiments in Fig. 20.9.4 jumps to 100% (since the fraction s/r of solved ones goes to zero).

Epistasis




Fig. 20.9.1: Unscaled ruggedness $\bar{st}$ plot.




Fig. 20.9.2: Scaled ruggedness $\bar{st}$ plot.




Fig. 20.9.3: Scaled deceptiveness – average success generations $\bar{st}$




Fig. 20.9.4: Scaled deceptiveness – failed runs

Fig. 20.9: Experiments with ruggedness and deceptiveness.




Fig. 20.10.1: Epistasis η and problem length n .




Fig. 20.10.2: Epistasis η and problem length n : failed runs.

Fig. 20.10: Experiments with epistasis.

Fig. 20.10.1 illustrates the relation between problem size n , the epistasis factor η , and the average success generations. Although rising epistasis makes the problems harder, the complexity does not rise as smoothly as in the previous experiments. The cause for this is likely the presence of crossover – if mutation was allowed solely, the impact of epistasis would most likely be more intense. Another interesting fact is that experimental settings with odd values of η tend to be much more complex than those with even ones. This relation becomes even more obvious in Fig. 20.10.2, where the proportion of failed runs, i. e., those which were not able to solve problem in less than 1000 generations, is plotted. A high plateau for greater values of η is cut by deep valleys at positions where $\eta = 2 + 2i \forall i \in \mathbb{N}$. This phenomenon has thoroughly been discussed by Niemczyk [1547] and can be excluded from the experiments by applying the scaling mechanism with parameter $y \in [0, 10]$ as defined in Equation 20.35:

$$\text{epistasis: } \eta = \begin{cases} y \leq 0 & \text{if } 2 \\ y \geq 10 & \text{if } 41 \\ 2 \lfloor 2y \rfloor + 1 & \text{otherwise} \end{cases} \quad (20.35)$$

Neutrality




Fig. 20.11: The results from experiments with neutrality.

Figure 20.11 illustrates the average number of generations $\overline{st}$ needed to grow an individual with optimal functional fitness for different values of the neutrality parameter μ . Until $\mu \approx 10$, the problem hardness increases rapidly. For larger degrees of redundancy, only minor increments in $\overline{st}$ can be observed. The reason for this strange behavior seems to be the crossover operation. Niemczyk [1547] shows that a lower crossover rate makes experiments involving the neutrality filter of the model problem very hard. We recommend using only μ -values in the range from zero to eleven for testing the capability of optimization methods of dealing with neutrality.

Epistasis and Neutrality

Our model problem consists of independent filters for the properties that may influence the hardness of an optimization task. It is especially interesting to find out whether these filters can be combined arbitrarily, i. e., if they are indeed free of interaction. In the ideal case, $\overline{st}$ of an experiment with $n = 80$, $\mu = 8$, and $\eta = 0$ added to $\overline{st}$ of an experiment for $n = 80$, $\mu = 0$, and $\eta = 4$ should roughly equal to $\overline{st}$ of an experiment with $n = 80$, $\mu = 8$, and $\eta = 4$. In Figure 20.12, we have sketched these expected values (Fig. 20.12.1) and the results of the corresponding real experiments (Fig. 20.12.2). In fact, these two diagrams are very similar. The small valleys caused by the “easier” values of η (see Section 20.2.6) occur in both charts. The only difference is a stronger influence of the degree of neutrality.




Fig. 20.12.1: The expected results.




Fig. 20.12.2: The experimentally obtained results.

Fig. 20.12: Expectation and reality: Experiments involving both, epistasis and neutrality

Ruggedness and Epistasis




Fig. 20.13.1: The expected results.




Fig. 20.13.2: The experimentally obtained results.

Fig. 20.13: Expectation and reality: Experiments involving both, ruggedness and epistasis

It is a well-known fact that epistasis leads to ruggedness, since it violates the causality as discussed in Section 1.4.5. Combining the ruggedness and the epistasis filter therefore leads to stronger interactions. In Fig. 20.13.2, the influence of ruggedness seems to be amplified by the presence of epistasis when compared with the estimated results shown in Fig. 20.13.1. Apart from this increase in problem hardness, the model problem behaves as expected. The characteristic valleys stemming from the epistasis filter are clearly visible, for instance.

Summary

In summary, this model problem has proven to be a viable approach for simulating problematic phenomena in optimization. It is

- functional, i. e., allows us to simulate many problematic features,
- tunable – each filter can be tuned independently,
- easy to understand,

- allows for very fast fitness computation,
- easily extensible – each filter can be replaced with other approaches for simulating the same feature.

Niemczyk [1547] has written a stand-alone Java class implementing the model problem which is provided at <http://www.sigoa.org/documents/> [accessed 2008-05-17] and <http://www.it-weise.de/documents/files/TunableModel.java> [accessed 2008-05-17]. This class allows setting the parameters discussed in this section and provides methods for determining the objective values of individuals in the form of `byte` arrays. In the future, some strange behaviors (like the irregularities in the ruggedness filter and the gaps in epistasis) of the model need to be revisited, explained, and, if possible, removed.

20.3 Genetic Programming Problems

20.3.1 Artificial Ant

We already have discussed parts of the Artificial Ant problem in Section 1.2.2 on page 9 – here we are going to investigate it more thoroughly. The goal of the original problem defined by Collins and Jefferson [1548], [1549, 1550] was to find a program that controls an artificial ant in a simulated environment. Such environments usually have the following features:

- It is divided in a toroidal grid generating rectangular cells in the plane making the positions of coordinates of all objects discrete.
- There exists exactly one ant in the environment.
- The ant will always be inside one cell at one time.
- A cell can either contain one piece of food or not.

The ant is a very simple life form. It always faces in one of the four directions north, east, south, or west. Furthermore, it can sense if there is food in the next cell in the direction it faces. It cannot sense if there is food on any other cell in the map.

Like space, the time in the Artificial Ant problem is also discrete. The ant may carry out one of the following actions per time unit:

- The ant can move for exactly one cell into the direction it faces. If this cell contains food, the ant consumes it in the very moment in which it enters the cell.
- The ant may turn left or right by 90.
- The ant may do nothing in a time unit.

Many researchers such as Collins and Jefferson [1548; 1550], [1549], Koza [65], Lee and Wong [222], Harries and Smith [1551], Luke and Spector [1552], Kuscü [1553], Chellapilla [1554], Ito et al. [1555], Langdon and Poli [1556], and Frey [1557] have ever since used the Artificial Ant problem as benchmark in their research. Since the Artificial Ant problem neither imposes a special genome, phenotype, nor otherwise restricts the parameters of the optimization process, it is the ideal environment for such tests. In order to make the benchmark results comparable, special instances of the problem like the *Santa Fe Trail* with well-defined features have been defined.

Santa Fe trail

One instance of the artificial ant problem is the “Santa Fe trail” sketched in Figure 20.14 designed by Langdon [65]. It is a map of 32 * 32 cells containing 89 food pellets distributed along a certain route. Initially, the ant will be placed in the upper left corner of the field facing east. In trail of food pellets, there are gaps of five forms:

- one cells along a straight line
- two cells along a straight line

- one cell in a corner
- two cells at a corner (requiring something like a “horse jump” in chess)
- three cells at a corner

The goal is here to find some form of control for the ant that allows it to eat as many of the food pellets as possible (the maximum is 89) and to walk a distance as short as possible in order to do so (the optimal route is illustrated in Figure 20.14). Of course, there will be a time limit set for the ant to perform this task (normally 200 time units).




Fig. 20.14: The Santa Fee Trail in the Artificial Ant Problem [65].

Solutions

Genetic Algorithm evolving Finite State Machines

Jefferson et al. [1549] applied a conventional genetic algorithm that evolved finite state machines encoded in a fixed-length binary string genome to the Artificial Ant problem. The sensor information together with the current state determines the next state, therefore a finite state machine with at most m states can be encoded in a chromosome using $2m$ genes. In order to understand the structure of such a chromosome, let us assume that $m = 2^n$. We then can specify the finite state machine as a table where $n + 1$ bits are used as row index. n of these indexes identify the current state and one bit is used for the sensor information (1=food ahead, 0=no food ahead). In total, there are $2m$ rows. There is no need to store the row indices, just the cell contents: n bits encode the next state, and two bits encode the action to be performed at the state transition (00 for nothing, 01 for turning left, 10 for

turning right, 11 for moving). A chromosome encoding a finite state machine with m states can be encoded in $2m(n+2) = 2^{n+1}(n+2)$ bits. If the initial state in the chromosome is also to be stored, another n bits are needed to do so. Every chromosome represents a valid finite state machine.

Jefferson et al. [1549] allowed for 32 states (453 bit chromosomes) in their finite state machines. They used one objective function that returned the number of food pellets eaten by the ant in a simulation run (of maximal 200 steps) and made it subject to maximization. Using a population of 65 536 individuals, they found one optimal solution (with fitness 89).

Genetic Algorithm evolving Artificial Neural Networks

Collins and Jefferson [1548] also evolved an artificial neural network (encoded in a 520 bit genome) with a genetic algorithm of the same population size to successfully solve the Artificial Ant problem. Later, they applied a similar approach [1550] with 25 590 bit genomes which allowed even the structure of the artificial neural networks to evolve to a generalized problem exploring the overall behavior of ant colonies from a more biological perspective.

Genetic Programming evolving Control Programs

Koza [1559; 801; 802] solved the Artificial Ant problem by evolving LISP¹¹-programs. Therefore, he introduced the parameterless instructions **MOVE**, **RIGHT**, and **LEFT** that moved the ant one unit, or turned it right or left respectively. Furthermore, the binary conditional expression **IF-FOOD-AHEAD** executed its first parameter expression if the ant could sense food and the second one otherwise. Two compound instructions, **PROGN2** and **PROGN3**, execute their two or three sub-expressions unconditionally. After 21 generations using a 500 individual population and fitness-proportional selection, Genetic Programming yielded an individual solving the Santa Fe trail optimally. Connolly [1560] provides an easy step-by-step guide for solving the Artificial Ant problem with Genetic Programming (and for visualizing this process).

20.3.2 The Greatest Common Divisor

Another problem suitable to test Genetic Programming approaches is to evolve an algorithm that computes the greatest common divisor¹², the GCD.

Problem Definition

Definition 20.1 (GCD). For two integer numbers $a, b \in \mathbb{N}_0$, the greatest common divisor (GCD) is the largest number c that divides both a ($c|a \equiv a \bmod c = 0$) and b ($c|b \equiv b \bmod c = 0$).

$$c = \gcd(a, b) \Leftrightarrow c|a \wedge c|b \wedge (\nexists d \in \mathbb{N} : d|a \wedge d|b \wedge d > c) \quad (20.36)$$

$$\Leftrightarrow \max \{e \in \mathbb{N} : (a \bmod e = 0) \wedge (b \bmod e = 0)\} \quad (20.37)$$

The Euclidean Algorithm

The GCD can be computed by the Euclidean algorithm¹³ which is specified in its original version in Algorithm 20.3 and in the improved, faster variant as Algorithm 20.4 [1561, 1562].

¹¹ http://en.wikipedia.org/wiki/Lisp_%28programming_language%29 [accessed 2007-07-03]

¹² http://en.wikipedia.org/wiki/Greatest_common_divisor [accessed 2007-10-05]

¹³ http://en.wikipedia.org/wiki/Euclidean_algorithm [accessed 2007-10-05]

Algorithm 20.3: $\text{gcd}(a, b) \leftarrow \text{euclidGcdOrig}(a, b)$

Input: $a, b \in \mathbb{N}_0$: two integers**Output:** $\text{gcd}(a, b)$: the greatest common divisor of a and b

```

1 begin
2   while  $b \neq 0$  do
3     if  $a > b$  then  $a \leftarrow a - b$ 
4     else  $b \leftarrow b - a$ 
5   return  $a$ 
6 end

```

Algorithm 20.4: $\text{gcd}(a, b) \leftarrow \text{euclidGcd}(a, b)$

Input: $a, b \in \mathbb{N}_0$: two integers**Data:** t : a temporary variable**Output:** $\text{gcd}(a, b)$: the greatest common divisor of a and b

```

1 begin
2   while  $b \neq 0$  do
3      $t \leftarrow b$ 
4      $b \leftarrow a \bmod b$ 
5      $a \leftarrow t$ 
6   return  $a$ 
7 end

```

The Objective Functions and the Prevalence Comparator

Although the GCD-problems seems to be more or less trivial since simple algorithms exist that solve it, it has characteristics that make it hard of Genetic Programming. Assume we have evolved a program $x \in \mathbb{X}$ which takes the two values a and b as input parameters and returns a new value $c = x(a, b)$. Unlike in symbolic regression¹⁴, it makes no sense to define the error between c and the real value $\text{gcd}(a, b)$ as objective function, since there is no relation between the “degree of correctness” of the algorithm and $|c - \text{gcd}(a, b)|$. Matter of fact, we cannot say that a program returning $c_1 = x_1(20, 15) = 6$ is better than $c_2 = x_2(20, 15) = 10$. 6 may be closer to the real result $\text{gcd}(20, 15) = 5$ but shares no divisor with it whereas $5 \mid 10 \equiv 10 \bmod 5 = 0$.

Based on the idea that the GCD of the variables a and b is preserved in each step of the Euclidean algorithm, a suitable functional objective function $f_1 : \mathbb{X} \mapsto [0, 5]$ for this problem is Algorithm 20.5. It takes a training case (a, b) as argument and first checks whether the evolved program $x \in \mathbb{X}$ returns the correct result for it. If so, $f_1(x) = 0$ is returned. Otherwise, we check if the greatest common divisor of $x(a, b)$ and a is still the greatest common divisor of a and b . If this is not the case, 1 is added to the objective value. The same is repeated with $x(a, b)$ and b . Furthermore, negative values of $x(a, b)$ are penalized with 2 and results that are larger or equal to a or b are penalized with one additional point for each violation. Depending on the program representation, this objective function is very rugged because small changes in the program have a large potential impact on the fitness. It also exhibits a large amount of neutrality, since it can take on only integer values between 0 (the optimum) and 5 (the worst case).

Additionally to f_1 , two objective functions optimizing non-functional aspects should be present. $f_2(x)$ minimizes the number of expressions in x and $f_3(x)$ minimizes the number of steps x needs until it terminates and returns the result. This way, we further small and fast algorithms. These three objective functions, combined to a prevalence¹⁵ comparator

¹⁴ See Section 22.1 for an extensive discussion of symbolic regression.¹⁵ See Definition 1.17 on page 20 for more information.

Algorithm 20.5: $r \leftarrow f_1(x, a, b)$

Input: $a, b \in \mathbb{N}_0$: the training case
Input: $x \in \mathbb{X}$: the evolved algorithm to be evaluated
Data: v : a variable holding the result of x for the training case
Output: r : the functional objective value of the functional objective function f_1 for the training case

```

1 begin
2    $r \leftarrow 0$ 
3    $v \leftarrow x(a, b)$ 
4   if  $v \neq \gcd(a, b)$  then
5      $r \leftarrow r + 1$ 
6     if  $\gcd(v, a) \neq \gcd(a, b)$  then  $r \leftarrow r + 1$ 
7     if  $\gcd(v, b) \neq \gcd(a, b)$  then  $r \leftarrow r + 1$ 
8     if  $v \leq 0$  then
9        $r \leftarrow r + 2$ 
10    else
11      if  $v \geq a$  then  $r \leftarrow r + 1$ 
12      if  $v \geq b$  then  $r \leftarrow r + 1$ 
13  return  $r$ 
14 end
  
```

$\text{cmp}_{F, \gcd}$, can serve as a benchmark on how good a Genetic Programming approach can cope with the rugged fitness landscape common to the evolution of real algorithms and how the parameter settings of the evolutionary algorithm influence this ability.

$$\text{cmp}_{F, \gcd}(x_1, x_2) = \begin{cases} -1 & \text{if } f_1(x_1) < f_1(x_2) \\ 1 & \text{if } f_1(x_1) > f_1(x_2) \\ \text{cmp}_{F, \text{Pareto}}(x_1, x_2) & \text{otherwise} \end{cases} \quad (20.38)$$

In principle, Equation 20.38 gives the functional fitness precedence before any other objective. If (and only if) the functional objective values of both individuals are equal, the prevalence is decided upon a Pareto comparison of the remaining two (non-functional) objectives.

The Training Cases

The structure of the training cases is also very important. If we simply use two random numbers a and b , their greatest common divisor is likely to 1 or 2. Hence, we construct a single training case by first drawing a random number $r \in \mathbb{N}$ uniformly distributed in $[10, 100\,000]$ as lower limit for the GCD and then keep drawing uniformly distributed random numbers $a > r, b > r$ until $\gcd(a, b) \geq r$. Furthermore, if multiple training cases are involved in the individual evaluation, we ensure that they involve different magnitudes of the values of a , b , and r . If we change the training cases after each generation of the evolutionary algorithm, the same goes for two subsequent training case sets. Some typical training sets are noted in Listing 20.2.

Generation 0		
=====		
a	b	gcd(a, b)
87546096	2012500485	21627
1656382	161406235	9101
7035	5628	1407
2008942236	579260484	972
556527320	1588840	144440
14328736	10746552	3582184

```

10 1390          268760      10
11 929436304    860551      5153
12 941094       1690414110  1386
13 14044248     1259211564  53604
14
15 Generation 1
16 =====
17 a            b            gcd(a,b)
18 117140       1194828      23428
19 2367         42080       263
20 3236545     379925       65
21 1796284190  979395390      10
22 4760        152346030   10
23 12037362    708102       186
24 1785869184  2093777664      61581696
25 782331      42435530     23707
26 434150199   24453828       63
27 45509100    7316463      35007
28
29 Generation 2
30 =====
31 a            b            gcd(a,b)
32 1749281113   82           41
33 25328611     99           11
34 279351072    2028016224   3627936
35 173078655    214140       53535
36 216          126          18
37 1607646156  583719700     2836
38 1059261     638524299   21
39 70903440    1035432       5256
40 26576383    19043        139
41 1349426     596258     31382

```

Listing 20.2: Some training cases for the GCD problem.

Rule-based Genetic Programming

We have conducted a rather large series of experiments on solving the GCD problem with Rule-based Genetic Programming (*RBGP*, see Section 4.8.3 on page 173). In this section, we will elaborate on the different parameters that we have tried out and what results could be observed for these settings.

Settings

As outlined in Table 20.4, we have tried to solve the GCD problem with Rule-based Genetic Programming with a lot of different settings (60 in total) which we will discuss here in accordance to Section 19.1.

Parameter	Short	Description
Problem Space	$\mathbb{X}$	The space of Rule-based Genetic Programming-programs with between 2 and 100 rules. (see Section 4.8.3)
Objective Functions	F	$F = \{f_1, f_2, f_3\}$ (see Section 20.3.2)

Search Space	$\mathbb{G}$	The variable-length bit strings with a gene size of 34 bits and a length between 64 and 3400 bits.
Search Operations	Op	30% mutation (including single-bit flits, permutations, gene deletion and insertion), 70% multi-point crossover
GPM	gpm	(see Figure 4.29)
Optimization Algorithm	alg	The optimization algorithm applied. $alg = 0 \rightarrow$ evolutionary algorithm $alg = 1 \rightarrow$ Parallel Random Walks
Comparison Operator	cm	(see Equation 20.38)
Population Size	ps	$ps \in \{512, 1024, 2048\}$
Steady-State	ss	The evolutionary algorithms were either steady state ($ss = 1$), meaning that the offspring had to compete with the already existing individuals in the selection phase, or generational/extinctive ($ss = 0$), meaning that only the offspring took part in the selection and the parents were discarded. (see Section 2.1.2) $ss = 0 \rightarrow$ generational $ss = 1 \rightarrow$ steady-state
Fitness Assignment Algorithm	fa	For fitness assignment in the evolutionary algorithm, Pareto ranking was used. (see Section 2.3.3)
Selection Algorithm	sel	A binary ($k = 2$) tournament selection was applied in the evolutionary algorithm. (see Section 2.4.4)
Convergence Prevention	cp	(see Section 2.4.8) $cp \in \{0, 0.3\}$
Number of Training Cases	tc	The number of training cases used for evaluating the objective functions. $tc \in \{1, 10\}$
Policy	ct	The policy according to which the training cases are changed. $ct = 0 \rightarrow$ The training cases do not change. $ct = 1 \rightarrow$ The training cases change each generation.
Generation Limit	mt	The maximum number of generations that each run is allowed to perform. (see Definition 1.43) $mt = 501$
System Configuration	Cfg	normal off-the-shelf PCs with approximately 2 GHz processor power

Table 20.4: The settings of the RBGP-Genetic Programming experiments for the GCD problem.

Convergence Prevention In our past experiments, we have made the experience that Genetic Programming in rugged fitness landscapes and Genetic Programming of real algorithms (which usually leads to rugged fitness landscapes) is very inclined to converge prematurely. If it finds some half-baked solution, the population often tended to converge to this individual and the evolutions stopped.

There are many ways to prevent this, like modifying the fitness assignment process by using sharing functions (see Section 2.3.4 on page 85), for example. Such methods influence individuals close in objective space and decrease their chance to reproduce. Here, we decided to choose a very simple measure which only decreases probability of reproduction of individuals with exactly equal objective functions: the simple convergence prevention algorithm introduced in Section 2.4.8. This filter has either been applied with strength $cp = 0.3$ or not been used ($cp = 0$).

Comparison with Random Walks We found it necessary to compare the Genetic Programming approach for solving this problem with random walks in order to find out whether or not

Genetic Programming can provide any advantage in a rugged fitness landscape. Therefore, we either used an evolutionary algorithm with the parameters discussed above ($alg = 0$) or parallel random walks ($alg = 1$). Random walks, in this context, are principally evolutionary algorithms where neither fitness assignment nor selection are preformed. Hence, we can test parameters like ps , ct , and tc , but no convergence prevention ($cp = 0$) and also no steady state ($expSteadyState = 0$). The results of these random walks are the best individuals encountered during their course.

Results

We have determined the following parameters from the data obtained with our experiments.

Measure	Short	Description
Perfection Fraction	p/r	The fraction of experimental runs that found perfect individuals. This fraction is also the estimate of the cumulative probability of finding a perfect individual until the 500th generation. (see Section 19.3.1)
Number of Perfect Runs	$\#p$	$p/r = CPp(ps, 500)$ (see Equation 19.20)
Number of Successful Runs	$\#s$	The number of runs where perfect individuals were discovered. (see Section 19.3.1)
Number of Comp. Runs	$\#r$	The number of runs where successful individuals were discovered. (see Section 19.3.1)
Mean Success Generation	$\overline{st}$	The total number of completed runs with the specified configuration.
Runs Needed for Perfection	pt_n	The average number of generations t needed by the (successful) experimental runs to find a successful individual. (or $\emptyset$ if no run was successful) (see Section 19.3.1)
Evaluations Needed for Perfection	$p\tau_n$	The estimated number $pt_n(0.99, ps, 500)$ of independent runs needed to find at least one perfect solution candidate with a probability of $z = 0.99$ until the 500th generation. (see Equation 19.21)
		The estimated number $pt_n(0.99, ps, 500)$ of objective function evaluations runs needed to find at least one perfect solution candidate with a probability of $z = 0.99$ until the 500th generation. (see Section 19.3.2)

Table 20.5: Evaluation parameters used in Table 20.6.

In the context of this experiment, a *perfect* solution represents a correct GCD algorithm, i.e., is not overfitted. Solutions with optimal functional objective values ($f_1 = 0$, whether due to overfitting or not) are called *successful*. Overfitted programs, like the one illustrated in Listing 20.4, will not work with inputs a and b different from those used in their evaluation.

Not all configurations were tested with the same number of runs since we had multiple computers involved in these test series and needed to end it at some point of time. We then used the maximum amount of available data for our evaluation. With the given configurations, the evolutionary algorithm runs usually took about one to ten minutes (depending on the population size). The results of the application of the Rule-based Genetic Programming to the GCD problem are listed in Table 20.6 below.

rank	alg	cp	ss	ct	tc	ps	p/r	$\#p$	$\#s$	$\#r$	$\overline{st}$	pt_n	$p\tau_n$
1.	0	0.3	1	0	1	1024	0.28	15	45	53	100.4	13.84	7 086 884
2.	0	0.3	1	0	1	512	0.12	6	35	51	98.5	36.79	9 419 095

3.	1	0	0	0	1	512	0.10	5	27	51	259.1	44.63	11 425 423
4.	0	0.3	1	0	10	2048	0.98	48	48	49	70.0	1.18	12 116 937
5.	1	0	0	0	1	1024	0.17	9	41	54	170.0	25.26	12 932 355
6.	0	0.3	0	0	1	2048	0.27	14	49	51	85.2	14.35	14 694 861
7.	0	0.3	1	0	10	1024	0.78	41	41	53	129.1	3.1	15 873 640
8.	0	0.3	1	0	1	2048	0.25	13	51	51	36.4	15.65	16 026 722
9.	0	0.3	1	0	10	512	0.49	25	25	51	153.0	6.84	17 498 481
10.	0	0.3	0	0	1	512	0.06	3	22	51	162.1	75.96	19 446 283
11.	0	0.3	0	1	10	1024	0.67	37	37	54	199.4	3.98	20 400 648
12.	0	0.3	0	1	1	1024	0.10	5	5	52	197.4	45.55	23 322 826
13.	0	0.3	1	1	10	2048	0.98	53	53	54	61.1	1.15	23 643 586
14.	0	0.3	0	0	1	1024	0.09	5	44	54	138.2	47.4	24 266 737
15.	0	0.3	0	0	10	2048	0.79	39	39	49	111.1	2.9	29 672 727
16.	0	0.3	0	1	10	2048	0.79	41	41	52	101.1	2.96	30 358 251
17.	0	0.3	1	1	10	1024	0.78	42	42	54	125.5	3.06	31 352 737
18.	0	0.3	1	1	1	2048	0.28	14	14	54	107.8	15.35	31 427 005
19.	0	0.3	0	0	10	1024	0.52	27	27	53	196.3	6.47	33 106 746
20.	0	0.3	0	1	10	512	0.25	13	13	52	231.5	16.01	40 980 085
21.	0	0.3	1	1	10	512	0.41	21	21	50	231.5	8.45	43 284 918
22.	0	0.3	0	1	1	2048	0.09	5	5	53	46.6	46.47	47 589 578
23.	0	0.3	0	0	10	512	0.19	10	10	52	250.8	21.56	55 199 744
24.	0	0.3	1	1	1	512	0.04	2	2	49	102.0	110.51	56 580 143
25.	0	0.3	1	1	1	1024	0.06	3	3	52	116.0	77.5	79 357 503
26.	0	0	0	0	10	1024	0.15	8	8	55	263.0	29.3	150 004 032
27.	0	0	1	0	10	1024	0.13	7	7	53	272.3	32.51	166 455 244
28.	0	0	0	0	10	2048	0.24	12	12	49	280.6	16.39	167 876 619
29.	0	0	1	0	1	2048	0.02	1	18	51	245.5	232.55	238 134 779
30.	1	0	0	0	1	2048	0.02	1	50	54	120.9	246.37	252 282 298
31.	0	0	1	0	10	2048	0.16	8	8	49	249.9	25.84	264 557 703
32.	0	0	1	1	10	512	0.08	4	4	50	320.3	55.23	282 777 841
33.	0	0	0	1	10	1024	0.06	3	3	53	264.3	79.03	404 649 274
34.	0	0	0	1	10	2048	0.10	5	5	50	237.4	43.71	447 576 992
35.	0	0	0	0	10	512	0.00	1	1	52	492.0	237.16	607 126 560
36.	0	0	1	1	10	2048	0.13	7	7	52	250.9	31.85	652 324 553
37.	0	0	1	1	10	1024	0.03	2	2	54	328.5	122.02	1 249 510 675
38.	1	0	0	1	1	2048	0.00	0	2	53	101.5	$+\infty$	$+\infty$
39.	0	0	0	0	1	2048	0.00	0	16	54	146.2	$+\infty$	$+\infty$
40.	0	0	1	0	1	512	0.00	0	6	51	202.0	$+\infty$	$+\infty$
41.	1	0	0	1	1	1024	0.00	0	2	53	209.0	$+\infty$	$+\infty$
42.	0	0	0	0	1	1024	0.00	0	9	54	257.1	$+\infty$	$+\infty$
43.	0	0	1	0	1	1024	0.00	0	16	54	277.3	$+\infty$	$+\infty$
44.	0	0	0	0	1	512	0.00	0	4	50	369.5	$+\infty$	$+\infty$
45.	0	0	0	1	1	1024	0.00	0	0	53	$\emptyset$	$+\infty$	$+\infty$

46.	0	0	0	1	1	2048	0.00	0	0	53	0	$+\infty$	$+\infty$
47.	0	0	0	1	1	512	0.00	0	0	51	0	$+\infty$	$+\infty$
48.	0	0	0	1	10	512	0.00	0	0	51	0	$+\infty$	$+\infty$
49.	0	0	1	0	10	512	0.00	0	0	52	0	$+\infty$	$+\infty$
50.	0	0	1	1	1	1024	0.00	0	0	52	0	$+\infty$	$+\infty$
51.	0	0	1	1	1	2048	0.00	0	0	54	0	$+\infty$	$+\infty$
52.	0	0	1	1	1	512	0.00	0	0	49	0	$+\infty$	$+\infty$
53.	0	0.3	0	1	1	512	0.00	0	0	52	0	$+\infty$	$+\infty$
54.	1	0	0	0	10	1024	0.00	0	0	55	0	$+\infty$	$+\infty$
55.	1	0	0	0	10	2048	0.00	0	0	49	0	$+\infty$	$+\infty$
56.	1	0	0	0	10	512	0.00	0	0	52	0	$+\infty$	$+\infty$
57.	1	0	0	1	1	512	0.00	0	0	51	0	$+\infty$	$+\infty$
58.	1	0	0	1	10	1024	0.00	0	0	53	0	$+\infty$	$+\infty$
59.	1	0	0	1	10	2048	0.00	0	0	51	0	$+\infty$	$+\infty$
60.	1	0	0	1	10	512	0.00	0	0	51	0	$+\infty$	$+\infty$

Table 20.6: Results of the RBGP test series on the GCD problem.

Each of the sixty rows of this table denotes one single test series. The first column contains the rank of the series according to $pt_n(., T, h)$ the following seven columns specify the settings of the test series as discussed in defined Table 20.4 on page 309. The last seven columns contain the evaluation results, which are formatted as follows:

Figure 20.15 illustrates the relation between the functional objective value f_1 of the currently best individual of the runs to the generation for the twelve best test series (according to their $p_{/r}$ -values). The curves are monotone for series with constant training sets ($ct = 0$) and jagged for those where the training data changed each generation ($ct = 1$).

```

1  false  ∨ true  ⇒ bt+1 = bt % at
2  (bt ≤ at) ∨ false ⇒ at+1 = at % bt
3  false  ∨ true  ⇒ ct+1 = bt

```

Listing 20.3: The RBGP version of the Euclidean algorithm.

```

1  (at ≤ bt)  ∧      true    ⇒ startt+1 = 1 - startt
2  false      ∨ (startt > at) ⇒ startt+1 = startt * 0
3  (at = 1)   ∧      (0 ≥ start) ⇒ startt+1 = startt / ct
4  true       ∧      (ct = startt) ⇒ ct+1 = ct + 1
5  (ct > 0)   ∨      (at ≤ bt) ⇒ at+1 = at * startt
6  true       ∧      true        ⇒ ct+1 = ct - ct
7  false      ∨      (at ≠ startt) ⇒ startt+1 = startt - startt
8  true       ∨      (ct = startt) ⇒ ct+1 = ct + 1
9  false      ∨      (0 < startt) ⇒ bt+1 = bt * ct
10 (startt = ct) ∨      (1 > startt) ⇒ bt+1 = bt % 1
11 (0 ≤ 1)     ∧      (0 ≥ 0)    ⇒ at+1 = at / ct
12 false      ∨      (bt < 0)   ⇒ at+1 = 1 - at
13 (startt ≤ startt) ∨      true    ⇒ ct+1 = ct / 0
14 (at = startt) ∧      true    ⇒ ct+1 = ct + 0
15 (at ≤ bt)   ∧      true    ⇒ startt+1 = 1 - startt

```

Listing 20.4: An overfitted RBGP solution to the GCD problem.




Fig. 20.15.1: $alg=0$,
 $cp=1$, $ss=1$, $ct=1$,
 $tc=10$, $ps=2048$




Fig. 20.15.2: $alg=0$,
 $cp=1$, $ss=1$, $ct=0$,
 $tc=10$, $ps=2048$




Fig. 20.15.3: $alg=0$,
 $cp=1$, $ss=0$, $ct=1$,
 $tc=10$, $ps=2048$




Fig. 20.15.4: $alg=0$,
 $cp=1$, $ss=0$, $ct=0$,
 $tc=10$, $ps=2048$




Fig. 20.15.5: $alg=0$,
 $cp=1$, $ss=1$, $ct=1$,
 $tc=10$, $ps=1024$




Fig. 20.15.6: $alg=0$,
 $cp=1$, $ss=1$, $ct=0$,
 $tc=10$, $ps=1024$




Fig. 20.15.7: $alg=0$,
 $cp=1$, $ss=0$, $ct=1$,
 $tc=10$, $ps=1024$




Fig. 20.15.8: $alg=0$,
 $cp=1$, $ss=0$, $ct=0$,
 $tc=10$, $ps=1024$




Fig. 20.15.9: $alg=0$,
 $cp=1$, $ss=1$, $ct=0$,
 $tc=10$, $ps=512$




Fig. 20.15.10: $alg=0$,
 $cp=1$, $ss=1$, $ct=1$,
 $tc=10$, $ps=512$




Fig. 20.15.11: $alg=0$,
 $cp=1$, $ss=1$, $ct=1$,
 $tc=1$, $ps=2048$




Fig. 20.15.12: $alg=0$,
 $cp=1$, $ss=1$, $ct=0$,
 $tc=1$, $ps=1024$

Fig. 20.15: The f_1 /generation-plots of the best configurations.

Discussion

We have sorted the runs in Table 20.6 according to their $p\tau_n$ -values, i.e., the (estimate of the) number of individual evaluations that are consumed by the pt_n independent experimental runs needed to find a perfect individual with $z = 99\%$ probability. Interestingly, our first approach was to evaluate these experiments according to their p/r -ratio, which led to a different order of the elements in the table.

Population Size The role of the population size ps is not obvious and there is no clear tendency which one to prefer when considering the $p\tau_n$ -value only. Amongst the three best EAs

according to this metric, we can find all three tested population sizes. When we focus on the p/r -ratio instead, the four best runs all have a population size of 2048. At least in this experiment, the bigger the population, the bigger the chance of success of an experiment holds. Of course, this comes with the trade-off that more individuals need to be processed which decreases $p\tau_n$. If we perform multiple runs with smaller populations, we have a high chance of finding at least one non-overfitted program with lesser objective function evaluations.

Steady State We can make a clear statement about the influence of the steady-state parameter ss had in our experiments. In almost all experimental runs, a configuration with $ss = 1$ was better than the exactly the same configuration with $ss = 0$ (compare, for instance, ranks 1 and 14 or rank 2 and 10 in Table 20.6). Also, if we look at the four best runs according to the p/r -rate, we can see that the better two of them both have $ss = 1$ while the other two have $ss = 0$ – while all other parameters remained constant. In this experiment, steady-state evolutionary algorithms outperform generational ones in both measures.

Convergence Prevention Even clearer is the influence of our primitive convergence prevention mechanism – the top 15 test series according to p/r all have $cp = 0.3$, and even generational tests with a population size of 512 beat steady-state runs with a population of 2048 individuals if using convergence prevention. Considering the estimated number $p\tau_n$ of individuals that need to be evaluated in pt_n independent runs to achieve 99% probability of finding a non-overfitted solution, this trend is even more obvious: amongst the 23 best Genetic Programming approaches all had the convergence prevention turned on. To be more precise: all but one single configuration with convergence prevention were better as all EA configurations with convergence prevention turned off. It seems that keeping the evolutionary process going and not allowing a single program to spread unchanged all throughout the population increases the solution quality a lot.

Number of Training Cases According to the $p\tau_n$ measure, using one training case $tc = 1$ is normally better than using $tc = 10$. In other words, we need fewer individual evaluations for finding a non-overfitted individual if fewer training cases are used. Obviously, using ten training cases corresponds to ten times as many individual evaluations per generation. When comparing row 1 and 7 in Table 20.6, the difference in estimated evaluations needed is only approximately two, and the configuration of row 23 needs approximately three times as many evaluations as row 10. The proportion of estimated extra evaluations is lower than the number of training cases used.

If we consider the fraction p/r of experiments that led to a perfect individual compared to the total number of experiments run for a configuration, the number of training cases has a very drastic effect: Then, the top ten test series all are based on ten training cases ($tc = 10$). We can think of a very simple reason for that which can be observed very well when comparing for example Fig. 20.15.12 with Fig. 20.15.9. In the twelfth best series, based on only a single training case ($tc = 1$) and illustrated in Fig. 20.15.12, only six values (0..5) for the objective function f_1 could occur. The ninth best series depicted in Fig. 20.15.9 on the other hand, had a much broader set of values of f_1 available. Since $tc = 10$ training cases were used and the final objective value assigned to an individual is the average of the scores reached in all these tests, it had much lower variations f_1 with $51 = |\{0.0, 0.1, 0.2, \dots, 4.8, 4.9, 5.0\}|$ levels. By using multiple training sets for these runs, we have effectively reduced the ruggedness of the fitness landscape and made it easier for the evolutionary algorithm to descend a gradient. The effect of increasing the resolution of the objective functions by increasing the number of training cases has also been reported in other researchers such as Lasarczyk and Banzhaf [1064] in the area of Algorithmic Chemistries¹⁶.

What we see is that a higher number of training cases significantly decreases overfitting and increases the chance of a run to find good solutions. It does, however, not decrease the expected number of individuals to be processed until a good solution is found.

¹⁶ You can find Algorithmic Chemistries discussed in Section 4.8.2 on page 171.

Changing Training Cases In this experiment, the EAs with constant training ($ct = 0$) cases clearly outperform those with training cases that change each generation ($ct = 1$) according to the $p\tau_n$ metric. This is strange, since one would expect that this approach would reduce overfitting and thus, since it does not a priori require more evaluations, improve the $p\tau_n$. This does also go for the speed of the evolution – the average first success generation $\overline{st}$ remains roughly constant, regardless if the training data changes or not. The best ten series according to $\overline{st}$ all use ten training cases ($tc = 10$), which seems to prevent overfitting sufficiently on its own. There is a difference in $tc = 1$, though, when we compare the perfect with runs with those which were just successful. In all runs that find a solution $x \in \mathbb{X}$ with $f_1(x) = 0$, this solution is also correct if $ct = 1$, i. e., $\#p = \#s$. In the test series where $ct = 0$, usually only a fraction of the runs that found an individual with optimal functional fitness had indeed found a solution. Here, overfitting takes place and $\#p < \#s$ can be usually observed.

In the context of this experiment, the parameter ct has no substantial positive influence on the chance of finding a solution to the GCD problem in a run. Instead, it even negatively influences the $p\tau_n$ values. Maybe the proportion of possible programs that are truly correct and those that just perform good on the test cases due to overfitting is relatively high in this problem. Then, the influence of this parameter could be different in other scenarios.

Comparison to Random Walks According to the chance p/r that a test series finds a non-overfitted solution, the best 17 configurations all were evolutionary algorithms, and apart from the 18th and 26th best series, no random walk made it into the top 30. Strangely, two random walks obtain very good placements (third and fifth rank) when considering the $p\tau_n$ metric but then, the next best random walk resides on rank 38. The two good random walks are configured in a way that leads to few evaluations, which leads to good values of $p\tau_n$ when accidentally a good solution was found. Nevertheless, having two of them in such high placements could either mean that the GCD problem is very hard (so searching with an EA is not really better than with a random walk) or very simple (since randomly created programs can solve it in many cases).

Be it how it be, the clear dominance of Genetic Programming in this problem indicates that there is a significant benefit of using it than just randomly picking individuals. Nevertheless, at least sometimes it seems to be a good idea to also try some runs which utilize the brute force of random walks when trying to solve a GP problem.

One of the reasons for the usual bad performance of the random walks was that the individuals tended to become unreasonable large. This also increased the amount of time needed for evaluation.

Contests

For most of the problems that can be solved with the aid of computers, a multitude of different approaches exist. They are often comparably good and their utility in single application cases strongly depends on parameter settings and thus, the experience of the user. Contests provide a stage where students, scientists, and practitioners from the industry can demonstrate their solutions to specific problems. They not only provide indications for which techniques are suitable for which tasks, but also give incitements and trickle scientific interest to improve and extend them. The RoboCup¹, for example is known to be the origin of many new, advanced techniques in robotics, image processing, cooperative behavior, multivariate data fusion, and motion controls [1563, 1564, 1565, 1566]². In this chapter, we discuss Genetic Programming approaches to competitions like the DATA-MINING-CUP or the Web Service Challenge.

21.1 DATA-MINING-CUP

21.1.1 Introduction

Data Mining

Definition 21.1 (Data Mining). Data mining³ can be defined as *the nontrivial extraction of implicit, previously unknown, and potentially useful information from data* [1567] and *the science of extracting useful information from large data sets or databases* [1568].

Today, gigantic amounts of data are collected in the web, in medical databases, by enterprise resource planning (ERP) and customer relationship management (CRM) systems in corporations, in web shops, by administrative and governmental bodies, and in science projects. These data sets are way too large to be incorporated directly into a decision making process or to be understood as-is by a human being. Instead, automated approaches have to be applied that extract the relevant information, to find underlying rules and patterns, or to detect time-dependent changes. Data mining subsumes the methods and techniques capable to perform this task. It is very closely related to estimation theory in stochastic (discussed in Section 27.6 on page 450) – the simplest summary of a data set is still the arithmetic mean of its elements. Data mining is also strongly related to artificial intelligence [1276, 1450], which includes learning algorithms that can generalize the given information. Some of the most wide spread and most common data mining techniques are:

¹ <http://www.robocup.org/> [accessed 2007-07-03] and <http://en.wikipedia.org/wiki/Robocup> [accessed 2007-07-03]

² Big up to the Carpe Noctem Robotic Soccer Team founded by my ingenious colleagues Baer and Reichle [1563] (<http://carpenoctem.das-lab.net/> [accessed 2008-04-23])!

³ http://en.wikipedia.org/wiki/Data_mining [accessed 2007-07-03]

- artificial neural networks (ANN) [1569, 1570],
- support vector machines (SVM) [1571, 1572, 1573, 1574],
- logistic regression [1575],
- decision trees [1576, 1577],
- Learning Classifier Systems as introduced in Chapter 7 on page 197, and
- naïve Bayes Classifiers [1578, 1579].

The DATA-MINING-CUP

The DATA-MINING-CUP⁴ (DMC) has been established in the year 2000 by the *prudsys AG*⁵ and the *Technical University of Chemnitz*⁶. It aims to provide an independent platform for data mining users and data analysis tool vendors and builds a bridge between academic science and economy. Today, it is one of Europe's biggest and most influential conferences in the area of data mining.

The DATA-MINING-CUP Contest is the biggest international student data mining competition. In the spring of each year, students of national and international universities challenge to find the best solution of a data analysis problem. Figure 21.1 shows the logos of the DMC from 2005 till 2007 obtained from <http://www.data-mining-cup.com/> [accessed 2007-07-03].



Fig. 21.1.1: 2005



Fig. 21.1.2: 2006



Fig. 21.1.3: 2007

Fig. 21.1: Some logos of the DATA-MINING-CUP.

21.1.2 The 2007 Contest – Using Classifier Systems

In Mai 2007, the students Stefan Achler, Martin Göb, and Christian Voigtmann came into my office and told me about the DMC. They knew that evolutionary algorithms are methods for global optimization that can be applied to a wide variety of tasks and wondered if they can be utilized for the DMC too. After some discussion about the problem to be solved, we together came up with the following approach which was then realized by them. While we are going to talk about our basic ideas and the results of the experiments, a detailed view on the implementation issues using the Java Sigoa framework are discussed in Section 25.1 on page 397. We have also summarized our work for this contest in a technical report [463].

⁴ The DATA-MINING-CUP is a registered trademark of prudsys AG. Der DATA-MINING-CUP ist eine eingetragene Marke der prudsys AG. <http://www.data-mining-cup.com/> [accessed 2007-07-03], <http://www.data-mining-cup.de/> [accessed 2007-07-03]

⁵ <http://www.prudsys.de/> [accessed 2007-07-03]

⁶ <http://www.tu-chemnitz.de> [accessed 2007-07-03] (Germany) – By the way, that's the university I've studied at, a great place with an excellent computer science department.

A Structured Approach to Data Mining

Whenever any sort of problem should be solved, a structured approach is always advisable. This goes for the application of optimization methods like evolutionary algorithms as well as for deriving classifiers in a data mining problem. In this section we discuss a few simple steps which should be valid for both kinds of tasks and which have been followed in our approach to the 2007 DMC.

The first step is always to clearly specify the problem that should be solved. Parts of this specification are possible target values and optimization criteria as well as the semantics of the problem domain. The optimization criteria tell us how different possible solutions can be compared with each other. If we were to sell tomatoes, for example, the target value (subject to maximization) would be the profit. Then again, the semantics of the problem domain allow us to draw conclusions on what features are important in the optimization or data mining process. When selling tomatoes, for instance, the average weight of the vegetables, their color, and maybe the time of the day when we open the store are important. The names of our customers on the other hand are probably not. The task of the DMC 2007 Contest, outlined in Section 21.1.2, is a good example for such a problem definition.

Before choosing or applying any data mining or optimization technique, an initial analysis of the given data should be performed. With this review and the problem specification, we can filter the data and maybe remove unnecessary features. Additionally, we will gain insight in the data structure and hopefully can already eliminate some possible solution approaches. It is, of course, better to exclude some mining techniques that cannot lead to good results in the initial phase instead of wasting working hours in trying them out to avail. Finding solutions with offline evolutionary computation usually takes a while, so we have now to decide on one or two solution approaches that are especially promising for the problem defined. We have performed this step for the DMC 2007 Contest data in Section 21.1.2 on page 321.

After this, we can apply the selected approaches. Of course, running an optimizer on all known sample data at once is not wise. Although we will obtain a result with which we can solve the specified problem for all the known data samples, it is possible not a good solution. Instead, it may be overfitted and can *only* process the data we were given. Normally however, we will only be provided with fraction of the “real data” and want to find a system that is able to perform well also on samples that are not yet known to us. Hence, we need to find out whether or not our approach generalizes. Therefore, solutions are derived for a subset of the available data samples only, the training data. These solutions are then tested on the test set, the remaining samples not used in its creations.⁷ The system we have created generalizes well if it is rated approximately equally good by the optimization criterion for both, the training and the test data. Now we can repeat the process by using all available data. We have evolved classifier systems that solve the DMC 2007 Contest according to this method in Section 21.1.2 on page 323.

The students Achler, Göb, and Voigtmann have participated in the 2007 DMC Contest and proceeded according to this pattern. In order to solve the challenge, they chose for a genetic algorithm evolving a fuzzy classifier system. The results of their participation are discussed in Section 21.1.2 on page 326. The following sub-sections are based on their experiences and impressions, and reproduce how they proceeded.

The Problem Definition

Rebate systems are an important means to animate customers to return to a store in classical retail. In the 2007 contest, we consider a check-out couponing system. Whenever a customer leaves a store, at the end of her bill a coupon can be attached. She then can use the coupon to receive some rebate on her next purchase. When printing the bill at the checkout, there are three options for couponing:

⁷ See also Section 1.4.6 for this approach.

Case N: attach no coupon to the bill,

Case A: attach coupon type **A**, a general rebate coupon, to the bill, or

Case B: attach coupon type **B**, a special voucher, to the bill.

The profit of the couponing system is defined as follows:

- Each coupon which is not redeemed costs 1 money unit.
- For each redeemed coupon of type **A**, the retailer gains 3 money units.
- For each coupon of type **B** which is redeemed, the retailer gains 6 money units.

It is thus clear that simply printing both coupons at the end of each bill makes no sense. In order to find a good strategy for coupon printing, the retailer has initiated a survey. She wants to find out which type of customer has an affinity to cash in coupons and, if so, which type of coupon most likely. Therefore the behavior of 50000 customers has been anonymously recorded. For all these customers, we know the customer ID, the number of redemptions of 20 different coupons and the historic information whether coupon type **A**, coupon type **B**, or none of them has been redeemed. Cases where both have been cashed in are omitted.

ID	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11	C12	C13	C14	C15	C16	C17	C18	C19	C20	Coupon
97006	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	N
97025	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	N
97032	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	A
97051	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	N
97054	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
97061	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	1	0	0	0	A
97068	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
97082	0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0	0	0	0	0	N
97093	0	0	0	1	0	0	1	0	0	0	0	1	0	0	1	0	1	0	0	0	B
97113	0	0	1	1	0	0	1	0	0	1	0	0	0	0	1	0	0	0	0	0	A
97128	1	1	0	0	0	0	1	0	0	1	0	0	0	0	1	0	0	0	0	0	N
97143	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	N
97178	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
97191	0	0	1	1	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
97204	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	N
97207	0	0	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	N
94101	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	N
94116	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
94118	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	N
94126	1	0	0	1	0	0	1	0	1	1	0	1	0	0	1	0	0	0	0	0	A
94129	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	N
94140	0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0	0	0	0	0	N
94143	0	0	0	1	0	0	1	0	0	1	0	0	0	0	1	0	0	0	0	0	N
94146	0	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	0	0	0	0	N
83151	0	1	0	0	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	N
83159	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	N
83162	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	N
83172	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	N
83185	0	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	N
83197	0	0	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	N
83203	0	0	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	N
83224	0	0	1	0	0	0	0	1	0	0	1	0	0	0	0	1	0	1	0	0	N
83229	0	0	0	0	1	0	0	0	0	1	0	0	0	0	1	0	1	0	0	0	N
83233	0	0	1	1	0	0	1	0	0	0	0	1	0	0	0	1	0	0	0	0	N
83235	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	N
83245	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	N
83259	0	0	0	1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	N
83264	0	0	1	0	0	0	0	1	0	0	1	0	0	0	1	0	0	0	0	0	N
83268	0	0	0	1	0	0	0	1	0	0	0	0	0	0	1	0	1	0	0	0	N
83276	0	1	1	1	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	N
83281	0	0	1	1	0	0	1	0	0	1	1	1	0	0	1	0	0	0	0	0	N
83285	0	0	1	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
83298	0	0	0	1	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	N
83315	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
83337	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	A
83347	0	0	0	1	0	0	1	0	0	0	0	0	0	0	1	0	0	0	0	0	N

Fig. 21.2: A few samples from the DMC 2007 training data.

Figure 21.2 shows some samples from this data set. The task is to use it as training data in order to derive a classifier C that is able to decide from a record of the 20 features whether a coupon **A**, **B**, or none should be provided to a customer. This means to maximize the profit $P(C)$ of retailer gained by using the classifier C which can be computed according to

$$P(C) = 3 * AA + 6 * BB - 1 * (NA + NB + BA + AB) \quad (21.1)$$

where

- AA is the number of correct assignments for coupon **A**.
- BB is the number of correct assignments for coupon **B**.
- NA is the number of wrong assignments to class **A** from the real class **N**.
- NB is the number of wrong assignments to class **B** from the real class **N**.
- BA is the number of wrong assignments to class **A** from the real class **B**.
- AB is the number of wrong assignments to class **B** from the real class **A**.

Wrong assignments from the classes **A** and **B** to **N** play no role.

The classifier built with the 50000 data samples is then to be applied to another 50000 data samples. There however, the column *Coupon* is missing and should be the result of the classification process. Based on the computed assignments, the profit score P is calculated for each contestant by the jury and the team with the highest profit will win.

Initial Data Analysis

The test dataset has some properties which make it especially hard for learning algorithms to find good solutions. Figure 21.3 for example shows three data samples with exactly the same features but different classes. In general, there is some degree of fuzzyness and noise, and clusters belonging to different classes overlap and contain each other. Since the classes

ID	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11	C12	C13	C14	C15	C16	C17	C18	C19	C20	Coupon
97054	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	N
94698	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	A
96366	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	B

Fig. 21.3: DMC 2007 sample data – same features but different classes.

cannot be separated by hyper-planes in a straightforward manner, the application of neural networks and support vector machines becomes difficult. Furthermore, the values of the features take on only four different values and are zero to 83.7%, as illustrated in Table 21.1. In general, such a small number of possible feature values makes it hard to apply methods

value	number of occurrences
0	837 119
1	161 936
2	924
3	21

Table 21.1: Feature-values in the 2007 DMC training sets.

that are based on distances or averages. Stefan, Martin, and Christian had already come to

this conclusion when we met. At least, one positive fact can easily be found by eyesight when inspecting the training data: the columns *C6*, *C14*, and *C20*, marked gray in Figure 21.2, are most probably insignificant since they are almost always zero and, hence, can be excluded from further analysis. The same goes for the first column, the customer *ID*, by common sense.

The Solution Approach: Classifier Systems

From the initial data analysis, we can reduce the space of values a feature may take on to 0, 1, and >1 . This limited, discrete range is especially suited for Learning Classifier Systems (LCS) discussed in Chapter 7 on page 197.

Since we already know the target function, $P(C)$, we do not need the learning part of the LCS. Instead, our idea was to use the profit $P(C)$ defined in Equation 21.1 directly as objective function for a genetic algorithm.

Very much like in the Pitt-approach [1069, 1261, 1070] in LCS, the genetic algorithm would be based on a population of classifier systems. Such a classifier system is a list of rules (the single classifiers). A rule contains a classification part and one condition for each feature in the input data. We used a two bit alphabet for the conditions, allowing us to encode the four different conditions per feature listed in Table 21.2. The three different classes can be

condition (in genotype)	condition (in phenotype)	corresponding feature value
00	0	must be 0
01	1	must be ≥ 1
10	2	must be > 1
11	3	do not care (i. e., any value is ok)

Table 21.2: Feature conditions in the rules.

represented using two additional bits, where 00 and 11 stands for **A**, 01 means **B**, and 10 corresponds to **N**. We leave three insignificant features away, so a rule is in total $17 \cdot 2 + 2 = 36$ bits small. This means that we need less memory for a classifier system with 17 rules than for 10 double precision floating point numbers, as used by a neural network, for example.

When a feature is to be classified, the rules of a classifier system are applied step by step. A rule fits to a given data sample if none of its conditions are violated by a corresponding sample feature. As soon as such a rule is found, the input is assigned to the class identified by the classification part of the rule. This stepwise interpretation creates a default hierarchy that allows classifications to include each other: a more specific rule (which is checked before the more general one) can represent a subset of features which is subsumed by a rule which is evaluated later. If no rule in the classifier systems fits to a data sample, **N** is returned per default since misclassifying an **A** or **B** as an **N** at least does not introduce a penalty in $P(C)$ according to Equation 21.1.

Since the input data is noisy, it turned out to be a good idea to introduce some fuzzyness in our classifiers, too, by modifying this default rule. During the classification process, we remember the rule which was violated by the least features. In the case that no rule fits perfectly, we check if the number of these misfits is less than one fifth of the features, in this case $\frac{17}{5} \approx 3$. If so, we consider it as a match and classify the input according to the rules classification part. Otherwise, the original default rule is applied and **N** is returned. Figure 21.4 outlines the relation of the genotype and phenotype of such a fuzzy classifier system. It shows a classifier system consisting of four rules that has been a real result of the genetic algorithm. In this graphic, we also apply it to the second sample of the dataset that is to be classified. As one can easily see, none of the four rules matches fully – which




Fig. 21.4: An example classifier for the 2007 DMC.

strangely is almost always the case for classifier systems that sprung of the artificial evolution performed by us. The data sample, however, violates only three conditions of the second rule and, hence, stays exactly at the $\frac{1}{5}$ -threshold. Since no other rule in the classifier system has less misfit conditions, the result of this classification process will be **A**.

Analysis of the Evolutionary Process

Table 21.3 lists the settings of the evolutionary algorithm that we have applied evolve classifiers for the DATA-MINING-CUP 2007 problem.

Parameter	Short	Description
Problem Space	$\mathbb{X}$	The space of classifiers consisting of between 2 and 55 rules. (see Section 21.1.2)
Objective Functions	F	$F = \{f_1, f_2\}$, where $f_1(C) = -P(C)$ rates the profit and $f_2(C) = \max\{\text{len}(C), 3\}$ is the non-functional length criterion.
Search Space	$\mathbb{G}$	The variable-length bit strings with a length between 74 and 2035 bits and a gene size of 37 bits. (see Section 3.5)
Search Operations	Op	70% multi-point crossover, 30% mutation (including single-bit flips, insertion, and deletion of genes)
GPM	gpm	(see Figure 21.4)
Optimization Algorithm	alg	elitist evolutionary algorithm (see ??)
Comparison	cm	Pareto comparison (see Section 1.2.4)
Operator Population Size	ps	$ps = 10\,243$

Maximum Archive Size	<i>as</i>	The size of the archive with the best known individuals was limited to $as = 101$. (see Definition 2.4)
Steady-State	<i>ss</i>	The algorithm was generational (not steady-state) ($ss = 0$). (see Section 2.1.2)
Fitness Assignment Algorithm	<i>fa</i>	For fitness assignment in the evolutionary algorithm, Pareto ranking was used. (see Section 2.3.3)
Selection Algorithm	<i>sel</i>	A binary ($k = 2$) tournament selection was applied. (see Section 2.4.4)
Convergence Prevention	<i>cp</i>	No additional means for convergence prevention were used, i. e., $cp = 0$. (see Section 2.4.8)
Generation Limit	<i>mt</i>	The maximum number of generations that each run is allowed to perform. (see Definition 1.43) $mt = 1001$

Table 21.3: The settings of the experiments for the DATA-MINING-CUP.




Fig. 21.5: The course of the classifier system evolution.

Figure 21.5 illustrates the course of the classifier system evolution. We can see a logarithmic growth of the profit with the generations as well as with the number of rules in the classifier systems. A profit of 8800 for the 50 000 data samples has been reached. Experiments with 10 000 datasets held back and an evolution on the remaining 40 000 samples indicated that the evolved rule sets generalize sufficiently well. The cause for the generalization of the results is the second, non-functional objective function which puts pressure into the direction of smaller classifier systems and the modified default rule which allows noisy input data. The result of the multi-objective optimization process is the Pareto-optimal set. It comprises all




Fig. 21.6: Some Pareto-optimal individuals among the evolved classifier systems.

solution candidates for which no other individual exists that is better in at least one objective value and not worse in any other. Figure 21.6 displays some classifier systems which are members of this set after generation 1000. *C1* is the smallest non-dominated classifier system. It consists of three rules which lead to a profit of 7222. *C2*, with one additional rule, reaches 7403. The 31-rule classifier system *C3* provides a gain of 8748 to which the system with the highest profit evolved, *C4*, adds only 45 to a total of 8793 with a trade-off of 18 additional rules (49 in total).

As shown in Table 21.1 on page 321, most feature values are 0 or 1 and there are only very few 2 and 3-valued features. In order to find out how different treatment of those will influence the performance of the classifiers and of the evolutionary process, we slightly modified the condition semantics in Table 21.4 by changing the meaning of rule 2 from > 1 to ≤ 1 (compare with Table 21.2 on page 322).

The progress of the evolution depicted in Figure 21.7 exhibits no significant difference to the first one illustrated in Figure 21.5. With the modified rule semantics, the best classifier

condition (in genotype)	condition (in phenotype)	corresponding feature value
00	0	must be 0
01	1	must be ≥ 1
10	2	must be ≤ 1
11	3	do not care (i. e., any value is ok)

Table 21.4: Different feature conditions in the rules.




Fig. 21.7: The course of the modified classifier system evolution.

system evolved delivered a profit of 8666 by utilizing 37 rules. This result is also not very much different from the original version. Hence, the treatment of the features with the values 2 and 3 does not seem to have much influence on the overall result. In the first approach, rule-condition 2 used them as distinctive criterion. The new method treats them the same as feature value 1, with slightly worse results.

Contest Results and Placement

A record number of 688 teams from 159 universities in 40 countries registered for the 2007 DMC Contest, from which only 248 were finally able to hand in results. The team of the RWTH Aachen won place one and two by scoring 7890 and 7832 points on the contest data set. Together with the team from the Darmstadt University of Technology, ranked third, they occupy the first eight placements. Our team reached place 29 which is quite a good result considering that none of its members had any prior experience in data mining.

Retrospectively, one can recognize that the winning gains are much lower than those we have discussed in the previous experiments. They are, however, results of the classification of a different data set – the profits in our experiment are obtained from the training sets and not from the contest data. Although our classifiers did generalize well in the initial tests, they seem to suffer from some degree of overfitting. Furthermore, the systems discussed

here are the result of reproduced experiments and not the original contribution from the students. The system with the highest profit that the students handed in also had gains around 8600 on the training sets. With a hill climbing optimizer, we squeezed out another 200, increasing, of course, the risk of additional overfitting. In the challenge, the best scored score of our team, a total profit of 7453 (5.5% less than the winning team). This classifier system was however grown with a much smaller population (4096) than in the experiments here, due to time restrictions.

It should also be noted that we did not achieve the best result with the best single classifier system evolved, but with a primitive combination of this system with another one: If both classifier systems delivered the same result for a record, this result was used. Otherwise, $\mathbf{N}$ was returned, which at least would not lead to additional costs (as follows from Equation 21.1 on page 321).

Conclusion

In order to solve the 2007 DATA-MINING-CUP contest we exercised a structured approach. After reviewing the data samples provided for the challenge, we have adapted the idea of classifier systems to the special needs of the competition. As a straightforward way of obtaining such systems, we have chosen a genetic algorithm with two objective functions. The first one maximized the utility of the classifiers by maximizing the profit function provided by the contest rules. The second objective function minimized a non-functional criterion, the number of rules in the classifiers. It was intended to restrict the amount of overfitting. The bred classifier systems showed reasonable good generalization properties on the test data sets separated from the original data samples, but seem to be overfitted when comparing these results with the profits gained in the contest. A conclusion is that it is hard to prevent overfitting in an evolution based on limited sample data – the best classifier system obtained will possibly be overfitted. In the challenge, the combination of two classifiers yielded the best results. Such combinations of multiple, independent systems will probably perform better than each of them alone.

In further projects, especially the last two conclusions drawn should be considered. Although we used a very simple way to combine our classifier systems for the contest, it still provided an advantage.

A classifier system in principle is nothing more but an estimator⁸. There exist many sophisticated methods of combining different estimators in order to achieve better results [1580]. The original version of such “boosting algorithms”, developed by Schapire [1581], theoretically allows to achieve an arbitrarily low error rate, requiring basic estimators with a performance only slightly better than random guessing on any input distribution. The AdaBoost algorithm by Freund and Schapire [1582; 1583] additionally takes into consideration the error rates of the estimators. With this approach, even classifiers of different architectures like a neural network and a Learning Classifier System can be combined. Since the classification task in the challenge required non-fuzzy answers in form of definite set memberships, the usage of weighted majority voting [1584, 1585], as already applied in a very primitive manner, would probably have been the best approach.

21.2 The Web Service Challenge

21.2.1 Introduction

Web Service Composition

The necessity for fast service composition systems and the overall idea of the WS-Challenge is directly connected with the emergence of Service-Oriented Architectures (SOA).

⁸ See our discussion on estimation theory in Section 27.6 on page 450.

Today, companies rely on IT-architectures which are as flexible as their business strategy. The software of an enterprise must be able to adapt to changes in the business processes like accounting, billing, the workflows, and even in the office software. If external vendors, suppliers, or customers change, the interfaces to their IT systems must be newly created or modified too. Hence, the architecture of corporate software has to be built with the anticipation of changes and updates [1586, 1587, 1588].

A SOA is the ideal architecture for such systems [1589, 1590]. Service oriented architectures allow us to modularize the business logic and to implement it in the form of services accessible in a network. Services are building blocks for service processes which represent the workflows of an enterprise. They can be added, removed, and updated at runtime without interfering with the ongoing business. A SOA can be seen as a complex system with manifold services as well as $n:m$ dependencies between services and applications:

- An application may need various service functionalities.
- Different applications may need the same service functionality.
- A certain functionality may be provided by multiple services.

Business now depends on the availability of service functionality, which is ensured by service management. Manual service management however becomes more and more cumbersome and ineffective with a rising number of relations between services and applications. Here, self-organization promises a solution for finding services that offer a specific functionality automatically.

Self-organizing approaches need a combination of syntactic and semantic service descriptions in order to decide whether a service provides a wanted functionality or not. Common syntactic definitions like WSDL [1591] specify the order and types of service parameters and return values. Semantic interface description languages like OWL-S [1592] or WSMO [1593, 1594] annotate these parameters with a meaning. While WSDL can be used to define a parameter `myisbn` of the type `String`, with OWL-S we can define that `myisbn` expects a `String` which actually contains an ISBN. Via a taxonomy we can now deduce that values which are annotated as either `ISBN-10` or `ISBN-13`⁹ can be passed to this service.

A wanted functionality is defined by a set of required output and available input parameters. A service offers this functionality if it can be executed with these available input parameters and its return values contain the needed output values. In order to find such services, the semantic concepts of their parameters are matched rather than their syntactic data types.

Many service management approaches employ semantic service discovery [1595, 1596, 1597, 1598, 1593, 1594, 1599, 1600]. Still, there is a substantial lack of research on algorithms and system design for fast response service discovery. This is especially the case in service composition where service functionality is not necessarily provided by a single service. Instead, combinations of services (*compositions*) are discovered. The sequential execution of these services provides the requested functionality.

The Web Service Challenge

Since 2005, the annual Web Service Challenge¹⁰ (WS-Challenge, WSC) provides a platform for researchers in the area of web service composition to compare their systems and exchange experiences [1601, 1602, 1603]. It is co-located with the IEEE Conference on Electronic Commerce (CEC) and the IEEE International Conference on e-Technology, e-Commerce, and e-Service (EEE).

Each team participating in this challenge provides one software system. A jury then uses these systems to solve different, complicated web service discovery and composition tasks.

⁹ There are two formats for International Standard Book Numbers (ISBNs), ISBN-10 and ISBN-13, see also <http://en.wikipedia.org/wiki/Isbn> [accessed 2007-09-02].

¹⁰ see <http://www.ws-challenge.org/> [accessed 2007-09-02]




Fig. 21.8: The logo of the Web Service Challenge.

The major evaluation criterion for the composers is the speed with which the problems are solved. Another criterion is the completeness of the solution. Additionally, there is also a prize for the best overall system architecture.

21.2.2 The 2006/2007 Semantic Challenge

We have participated in the 2006 and 2007 Web Service Challenges [1604, 1605]. Here we present the system, algorithms and data structures for semantic web service composition that we applied in both challenges. A slightly more thorough discussion of this topic can be found in Weise et al. [1606; 1607]. The tasks of the 2006 Web Service Challenge in San Francisco, USA and the 2007 WSC in Tokyo, Japan are quite similar and only deviate in the way in which the solutions have to be provided by the software systems. Hence, we will discuss the two challenges together in this single section. Furthermore, we only consider the semantic challenges, since they are more demanding than mere syntactic matching.

Semantic Service Composition

In order to discuss the idea of semantic service composition properly, we need some prerequisites. Therefore, let us initially define the set of all semantic concepts $\mathbb{M}$. All concepts that exist in the knowledge base are members of $\mathbb{M}$ and can be represented as nodes in a wood of taxonomy trees.

Definition 21.2 (subsumes). Two concepts $A, B \in \mathbb{M}$ can be related in one of four possible ways. We define the predicate $\text{subsumes} : \mathbb{M} \times \mathbb{M} \mapsto \mathbb{B}$ to express this relation as follows:

1. $\text{subsumes}(A, B)$ holds if and only if A is a generalization of B (B is then a specialization of A).
2. $\text{subsumes}(B, A)$ holds if and only if A is a specialization of B (B is then a generalization of A).
3. If neither $\text{subsumes}(A, B)$ nor $\text{subsumes}(B, A)$ holds, A and B are not related to each other.
4. $\text{subsumes}(A, B)$ and $\text{subsumes}(B, A)$ is **true** if and only if $A = B$ (antisymmetrie, as defined in Equation 26.59 on page 415).

The subsumes relation is transitive (see Equation 26.55 on page 413), and so are generalization and specialization: If A is a generalization of B ($\text{subsumes}(A, B)$) and B is a generalization of C ($\text{subsumes}(B, C)$), then A is also a generalization of C ($\text{subsumes}(A, C)$). The same goes vice versa for specialization, here we can define that if A is a specialization of B ($\text{subsumes}(B, A)$) and A is also a specialization of C ($\text{subsumes}(C, A)$), then either $\text{subsumes}(B, C)$ or $\text{subsumes}(C, B)$ (or both) must hold, i.e., either C is a specialization of B , or B is a specialization of C , or $B = C$.

If a parameter x of a service is annotated with A and a value y annotated with B is available, we can set $x = y$ and call the service only if $\text{subsumes}(A, B)$ holds (*contravariance*). This means that x expects less or equal information than given in y . The hierarchy defined here is pretty much the same as in object-oriented programming languages. If we imagine **A** and **B** to be classes in Java, $\text{subsumes}(A, B)$ can be considered to be equivalent to the expression `A.class.isAssignableFrom(B.class)`. If it evaluates to `true`, a value y of type **B** can be assigned to a variable x of type **A** since `y instanceof A` will also be `true`.

From the viewpoint of a composition algorithm, there is no need for a distinction between parameters and the annotated concepts. The set $\mathbb{S}$ contains all the services s known to the service registry. Each service $s \in \mathbb{S}$ has a set of required input concepts $s.in \subseteq \mathbb{M}$ and a set of output concepts $s.out \subseteq \mathbb{M}$ which it will deliver on return. We can trigger a service if we can provide all of its input parameters.

Similarly, a composition request R always consists of a set of available input concepts $R.in \subseteq \mathbb{M}$ and a set of requested output concepts $R.out \subseteq \mathbb{M}$. A composition algorithm in the sense of the Web Service Challenges 2006 and 2007 discovers a (topologically sorted¹¹) set of n services $S = \{s_1, s_2, \dots, s_n\} : s_1, \dots, s_n \in \mathbb{S}$. As shown in Equation 21.2, the first service (s_0) of a valid composition can be executed with instances of the input concepts $R.in$. Together with $R.in$, its outputs ($s_1.out$) are available for executing the next service (s_2) in S , and so on. The composition provides outputs that are either annotated with exactly the requested concepts $R.out$ or with more specific ones (*covariance*). Assuming that $R.in \cap R.out = \emptyset$, for each composition solving the request R , the predicate $\text{isGoal}(S)$ will hold. With Equation 21.2, we have defined the goal predicate which we can use in any form of informed or uninformed state space search (see Chapter 16 on page 245).

$$\begin{aligned} \text{isGoal}(S) \Leftrightarrow & \forall A \in s_1.in \exists B \in R.in : \text{subsumes}(A, B) \wedge \\ & \forall A \in s_i.in, i \in \{2..n\} \exists B \in R.in \cup s_{i-1}.out \cup \dots \cup s_1.out : \text{subsumes}(A, B) \wedge \\ & \forall A \in R.out \exists B \in s_1.out \cup \dots \cup s_n.out : \text{subsumes}(A, B) \end{aligned} \quad (21.2)$$

The Problem Definition

In the 2006 and 2007 Web Service Challenge, the composition software is provided with three parameters:

1. A concept taxonomy to be loaded into the knowledge base of the system. This taxonomy was stored in a file of the XML Schema format [1608].
2. A directory containing the specifications of the service to be loaded into the service registry. For each service, there was a single file given in WSDL format [1591].
3. A query file containing multiple service composition requests $R_1, R_2, \dots$ in a made-up XML [1609] format.

These formats are very common and allow the contestants to apply the solutions in real world applications later as well as using customized versions of their already existing applications. The expected result to be returned by the software was also a stream of data in a proprietary XML dialect containing all possible service compositions that solved the

¹¹ The set S is only partially ordered since, in principle, some services may be executed in parallel if they do not depend on each other.

queries according to Equation 21.2. It was possible that a request R_i was resolved by multiple service compositions. In the 2006 challenge, the communication between the jury and the programs was via command line or other interfaces provided by the software, in 2007 a web service interface was obligatory.

We will not discuss the data formats used in this challenge any further since they are replaceable and do not contribute to the way the composition queries are solved. Remarkably, however, were the following restrictions in the challenge tasks:

- There exists at least one solution for each query.
- The services in the solutions are represented as a sequence of sets. Each set contains equivalent services. Executing one service from each set forms a valid composition S . This representation does not allow for any notation of parallelization.

Before we elaborate on the solution itself, let us define the operation “promising” which obtains the set of all services $s \in \mathbb{S}$ that produce an output parameter annotated with the concept A (regardless of their inputs).

$$\forall A \in \mathbb{S}, \forall s \in \text{promising}(A) \Rightarrow \exists B \in s.out : \text{subsumes}(A, B) \quad (21.3)$$

The composition system that we have applied in the 2007 WSC consists of three types of composition algorithms. The problem space $\mathbb{X}$ that they investigate is basically the set of all possible permutations of all possible sets of services. The power set $\mathcal{P}(\mathbb{S})$ includes all possible subsets of $\mathbb{S}$. $\mathbb{X}$ is the set of all possible permutations of the elements in such subsets, in other words $\mathbb{X} \subseteq \{\forall \text{ permutation}(\xi) : \xi \in \mathcal{P}(\mathbb{S})\}$.

An (Uninformed) Algorithm Based on IDDFS

The most general and straightforward approach to web service composition is the uninformed search, an iterative deepening depth-first search (IDDFS) algorithm as discussed in Section 16.3.4 on page 249. Uninformed search algorithms do not make use of any information different from goal predicates as defined in Equation 21.2. We can build such a composition algorithm based on iterative deepening depth-first search. It is only fast in finding solutions for small service repositories but optimal if the problem requires an exhaustive search. Thus, it may be used by the strategic planner in conjunction with another algorithm that runs in parallel if the size of the repository is reasonable small or if it is unclear whether the problem can actually be solved.

Algorithm 21.1 (`webServiceCompositionIDDFS`) builds a valid web service composition starting from the back. In each recursion, its internal helper method `dl_dfs_wsc` tests all elements A of the set *wanted* of yet unknown parameters. It then iterates over the set of all services s that can provide A . For every single s , *wanted* is recomputed. If it becomes the empty set $\emptyset$, we have found a valid composition and can return it. If `dl_dfs_wsc` is not able to find a solution within the maximum depth limit (which denotes the maximum number of services in the composition), it returns $\emptyset$. The loop in Algorithm 21.1 iteratively invokes `dl_dfs_wsc` by increasing the depth limit step by step, until a valid solution is found.

An (Informed) Heuristic Approach

The IDDFS-algorithm just discussed performs an uninformed search in the space of possible service compositions. As we know from Section 16.4 on page 251, we can increase the search speed by defining good heuristics and using domain information. Such information can easily be derived in this research area. Therefore, we will again need some further definitions. Notice that the set functions specified in the following does not need to be evaluated every time they are queried, since we can maintain their information as meta-data along with the composition and thus save runtime.

Algorithm 21.1: $S \leftarrow \text{webServiceCompositionIDDFS}(R)$

Input: R : the composition request
Data: $\text{maxDepth}, \text{depth}$: the maximum and the current search depth
Data: in, out : current parameter sets
Output: S : a valid service composition solving R

```

1 begin
2    $\text{maxDepth} \leftarrow 2$ 
3   repeat
4      $S \leftarrow \text{dl\_dfs\_wsc}(R.\text{in}, R.\text{out}, \emptyset, 1)$ 
5      $\text{maxDepth} \leftarrow \text{maxDepth} + 1$ 
6   until  $S \neq \emptyset$ 
7   Subalgorithm  $S \leftarrow \text{dl\_dfs\_wsc}(\text{in}, \text{out}, \text{composition}, \text{depth})$ 
8   begin
9     foreach  $A \in \text{out}$  do
10      foreach  $s \in \text{promising}(A)$  do
11         $\text{wanted} \leftarrow \text{out}$ 
12        foreach  $B \in \text{wanted}$  do
13           $\text{if } \exists C \in s.\text{out} : \text{subsumes}(B, C)$  then  $\text{wanted} \leftarrow \text{wanted} \setminus \{B\}$ 
14          foreach  $D \in s.\text{in}$  do
15             $\text{if } \nexists E \in \text{in} : \text{subsumes}(D, E)$  then  $\text{wanted} \leftarrow \text{wanted} \cup \{D\}$ 
16           $\text{comp} \leftarrow s \oplus \text{composition}$ 
17          if  $\text{wanted} = \emptyset$  then
18            return  $\text{comp}$ 
19          else
20            if  $\text{depth} < \text{maxDepth}$  then
21               $\text{comp} \leftarrow \text{dl\_dfs\_wsc}(\text{in}, \text{wanted}, \text{comp}, \text{depth} + 1)$ 
22            if  $\text{comp} \neq \emptyset$  then return  $\text{comp}$ 
23      return  $\emptyset$ 
24   end
25 end

```

Let us first define the set of unsatisfied parameters $\text{wanted}(S) \subseteq \mathbb{M}$ in a candidate composition S as

$$\text{wanted}(S) \Leftrightarrow \{A : s_i \in S \wedge A \in s_i.\text{in} \wedge A \notin R.\text{in} \wedge (\nexists s_j \in S : 0 \leq j < i \wedge A \in s_j.\text{out})\} \cup R.\text{out} \setminus (R.\text{in} \cup \bigcup_{s \in S} s.\text{out}) \quad (21.4)$$

In other words, a wanted parameter is either an output concept of the composition query or an input concept of any of the services in the composition candidate that has not been satisfied by neither an input parameter of the query nor by an output parameter of any service. Here we assume that the concept A wanted by service s is not also an output parameter of s . This is done for simplification purposes – the implementation has to keep track of this possibility.

The set of eliminated parameters of a service composition contains all input parameters of the services of the composition and queried output parameters of the composition request that already have been satisfied.

$$\text{eliminated}(S) = \left(R.\text{out} \cup \bigcup_{s \in S} s.\text{in} \right) \setminus \text{wanted}(S) \quad (21.5)$$

Finally, the set of known concepts is the union of the input parameters defined in the composition request and the output parameters of all services in the composition candidate.

$$\text{known}(S) = R.in \cup \bigcup_{\forall s \in S} s.out \quad (21.6)$$

Instead of using these sets to build a heuristic function h , we can derive a comparator function cmp_{wsc} directly. This comparator function has the advantage that we also can apply randomized optimization methods like evolutionary algorithms based on it.

Algorithm 21.2: $r \leftarrow \text{cmp}_{wsc}(S_1, S_2)$

Input: S_1, S_2 : two composition candidates

Output: $r \in \mathbb{Z}$: indicating whether S_1 ($r < 0$) or S_2 ($r > 0$) should be expanded next

```

1 begin
2    $i_1 \leftarrow \text{len}(\text{wanted}(S_1))$ 
3    $i_2 \leftarrow \text{len}(\text{wanted}(S_2))$ 
4   if  $i_1 = 0$  then
5     if  $i_2 = 0$  then return  $\text{len}(S_1) - \text{len}(S_2)$ 
6     else return  $-1$ 
7   if  $i_2 = 0$  then return  $1$ 
8    $e_1 \leftarrow \text{len}(\text{eliminated}(S_1))$ 
9    $e_2 \leftarrow \text{len}(\text{eliminated}(S_2))$ 
10  if  $e_1 > e_2$  then return  $1$ 
11  else if  $e_1 < e_2$  then return  $-1$ 
12  if  $i_1 < i_2$  then return  $-1$ 
13  else if  $i_1 > i_2$  then return  $1$ 
14  if  $\text{len}(S_1) \neq \text{len}(S_2)$  then return  $\text{len}(S_1) - \text{len}(S_2)$ 
15  return  $\text{len}(\text{known}(S_1)) - \text{len}(\text{known}(S_2))$ 
16 end
```

Algorithm 21.2 defines cmp_{wsc} which compares two composition candidates S_1 and S_2 . This function can be used by a greedy search algorithm in order to decide which of the two possible solutions is more prospective. cmp_{wsc} will return a negative value if S_1 seems to be closer to a solution than S_2 , a positive value if S_2 looks as if it should be examined before S_1 , and zero if both seem to be equally good.

First, it compares the number of wanted parameters. If a composition has no such unsatisfied concepts, it is a valid solution. If both, S_1 and S_2 are valid, the solution involving fewer services wins. If only one of them is complete, it also wins. Otherwise, both candidates still have unsatisfied concepts. Only if both of them have the same number of satisfied parameters, we again compare the wanted concepts. For us, it was surprising that using the number of already satisfied concepts as comparison criterion with a higher priority than the number of remaining unsatisfied concepts. However, if we do so, the search algorithms perform significantly faster. If their numbers are also equal, we prefer the shorter composition candidate. If even the compositions are of the same length, we finally base the decision on the total number of known concepts. The interesting form of this comparator function is maybe caused by the special requirements of the WSC data. Nevertheless, it shows which sorts of information about a composition can be incorporated into the search.

Using such the comparator function cmp_{wsc} , we can customize the greedy search approach defined in Algorithm 16.6 on page 252 for web service composition. The function `greedyComposition` defined in Algorithm 21.3 performs such a greedy composing by maintaining an internal list which is descendingly sorted according to cmp_{wsc} . In each iteration, the last element is popped from the list and either returned (if it is a valid composition) or expanded by appending services providing wanted concepts.

Algorithm 21.3: $S \leftarrow \text{greedyComposition}(R)$

Input: R : the composition request
Data: X : the descendingly sorted list of compositions to explore
Output: S : the solution composition found, or $\emptyset$

```

1 begin
2    $X \leftarrow \bigcup_{A \in R.out} \text{promising}(A)$ 
3   while  $X \neq \emptyset$  do
4      $X \leftarrow \text{sortList}_d(X, \text{wanted})$ 
5      $S \leftarrow X_{[\text{len}(X)-1]}$ 
6      $X \leftarrow \text{deleteListItem}(X, \text{len}(X) - 1)$ 
7     if  $\text{isGoal}(S)$  then
8       return  $S$ 
9     foreach  $A \in \text{wanted}(S)$  do
10      foreach  $s \in \text{promising}(A)$  do
11         $X \leftarrow \text{addListItem}(X, s \oplus S)$ 
12  return  $\emptyset$ 
13 end

```

An Evolutionary Approach

In order to use a evolutionary algorithm to breed web service compositions, we first need to define a proper genome $\mathbb{G}$ able to represent service sequences. A straightforward yet efficient way is to use (variable-length) strings of service identifiers which can be processed by standard genetic algorithms (see Section 3.5 on page 121). A service can be identified by a number from $\mathbb{N}_0$ denoting its index in the list of all services in the registry. The genotype-phenotype mapping transforming the genotypes $g \in \mathbb{G}$ which are sequences of such identifiers to sequences of services, i. e., the phenotypes $S \in \mathbb{X}$, is thus trivial.

Because of the well-known string form, we could apply the standard creation, mutation, and crossover operators. However, by specifying a specialized mutation operation, we can make the search more efficient. This new operation either deletes the first service in S (via mutate_{wsc1}) or adds a promising service to S (as done in mutate_{wsc2}). Using the adjustable variable σ as a threshold we can tell the search whether it should prefer growing or shrinking the solution candidates.

$$\text{mutate}_{wsc1}(S) \equiv \begin{cases} \{s_2, s_3, \dots, s_{\text{len}(S)}\} & \text{if } \text{len}(S) > 1 \\ S & \text{otherwise} \end{cases} \quad (21.7)$$

$$\text{mutate}_{wsc2}(S) \equiv s \oplus S : A \in \text{wanted}(S) \wedge s \in \text{promising}(A) \quad (21.8)$$

$$\text{mutate}_{wsc}(S) \equiv \begin{cases} \text{mutate}_{wsc1}(S) & \text{if } \text{random}_u() > \sigma \\ \text{mutate}_{wsc2}(S) & \text{otherwise} \end{cases} \quad (21.9)$$

A new create operation for building the initial random configurations can be defined as a sequence of mutate_{wsc2} invocations of random length. Initially, $\text{mutate}_{wsc2}(\emptyset)$ will return a composition consisting of a single service that satisfies at least one parameter in $R.out$. We iteratively apply mutate_{wsc2} to its previous result a random number of times in order to create a new individual.

The Comparator Function and Pareto Optimization

As driving force for the evolutionary process we can reuse the comparator function cmp_{wsc} as specified as for the greedy search in Algorithm 21.2 on the previous page. It combines multiple objectives, putting pressure towards the direction of

- compositions which are complete,

- small compositions,
- compositions that resolve many unknown parameters, and
- compositions that provide many parameters.

On the other hand, we could as well separate these single aspects into different objective functions and apply direct Pareto optimization. This has the drawback that it spreads the pressure of the optimization process over the complete Pareto frontier¹².




Fig. 21.9: A sketch of the Pareto front in the genetic composition algorithm.

Figure 21.9 visualizes the multi-objective version of the optimization problem “web service composition” by sketching a characteristic example for Pareto frontiers of several generations of an evolutionary algorithm. We concentrate on the two dimensions *composition size* and *number of wanted (unsatisfied) parameters*. Obviously, we need to find compositions which are correct, i. e., where the latter objective is zero. On the other hand, an evolution guided only by this objective can (and will) produce compositions containing additional, useless invocations of services not related to the problem at all. The size objective is thus also required.

In the manufactured example depicted in Figure 21.9, the first five or so generations are not able to produce good compositions yet. We just can observe that longer compositions tend to provide more parameters (and have thus a lower number of wanted parameters). In generation 20, the Pareto frontier is pushed farther forward and touches the abscissa – the first correct solution is found. In the generations to come, this solution is improved and useless service calls are successively removed, so the composition size decreases. There will be a limit, illustrated as generation 50, where the shortest compositions for all possible values of wanted are found. From now on, the Pareto front cannot progress any further and the optimization process has come to a rest.

As you can see, pure Pareto optimization does not only seek for the best correct solution but also looks for the best possible composition consisting of only one service, for the best one with two service, with three services, and so on. This spreading of the population of course slows down the progress into the specific direction where $wanted(S)$ decreases.

The comparator function cmp_{wsc} proven to be more efficient in focusing the evolution on this part of the problem space. The genetic algorithm based on it is superior in performance and hence, is used in our experiments.

¹² See Section 1.2.4 on page 15 for a detailed discussion on the drawbacks of pure Pareto optimization.

Experimental Results

Test	Depth of Solution	No. of Concepts	No. of Services	IDDFS (ms)	Greedy (ms)	GA (ms)
1	5	56 210	1000	241	34	376
2	12	56 210	1000	-	51	1011
3	10	58 254	10 000	-	46	1069
4	15	58 254	2000	-	36	974
5	30	58 254	4000	-	70	6870
6	40	58 254	8000	-	63	24 117
7	1	1590	118	≤16	≤16	290
8.1	2	15 540	4480	≤16	≤16	164
8.2	2	15 540	4480	≤16	≤16	164
8.3	2	15 540	4480	≤16	≤16	164
8.4	2	15 540	4480	≤16	≤16	234
8.5	3	15 540	4480	≤16	≤16	224
8.6	3	15 540	4480	≤16	≤16	297
8.7	4	15 540	4480	18	24	283
8.8	3	15 540	4480	≤16	≤16	229
8.9	2	15 540	4480	≤16	≤16	167
11.1	8	10 890	4000	-	31	625
11.3	2	10 890	4000	-	21	167
11.5	4	10 890	4000	22 021	≤16	281
12.1	5	43 680	2000	200 320	≤16	500
12.3	7	43 680	2000	99	31	375
13	6	43 680	2000	250	32	422

Table 21.5: Experimental results for the web service composers.

In Table 21.5, we illustrate the times that the algorithms introduced in this section needed to perform composition tasks of different complexity¹³. We have repeated the experiments multiple times on an off-the-shelf PC¹⁴ and noted the mean values. The times themselves are not so important, rather are the proportions and relations between them.

The IDDFS approach can only solve smaller problems and becomes infeasible very fast. When building simpler compositions though, it is about as fast as the heuristic approach, which was clearly dominating in all categories. A heuristic may be misleading and (although this didn't happen in our experiments) could lead to a very long computation time in the worst case. Furthermore, if a problem cannot be solved, the heuristic will not be faster than an uninformed search. Thus, we decided to keep both, the IDDFS and the heuristic approach in our system and run them in parallel on each task if sufficient CPUs are available.

The genetic algorithm (population size 1024, tournament selection) was able to resolve all composition requests correctly for all knowledge bases and all registry sizes. It was able to build good solutions regardless how many services had to be involved in a valid solution (solution depth). In spite of this correctness, it always was a magnitude slower than the greedy search which provided the same level of correctness.

If the compositions would become more complicated or involve quality of service (QoS) aspects, it is not clear if these can be resolved with a simple heuristic. Then, the genetic algorithm could outperform greedy search approaches.

¹³ The test sets used here are available at http://www.it-weise.de/documents/files/BWG2007WSC_software.zip [accessed August 12, 2008]. Well, at least partly, I've accidentally deleted set 12 and 13. Sorry.

¹⁴ 2GHz, Pentium IV single core with Hyper-Threading, 1GiB RAM, Windows XP, Java 1.6.0._03-b05

Architectural Considerations

In 2007, we introduced a more refined version [1605] of our 2006 semantic composition system [1604]. The architecture of this composer, as illustrated in Figure 21.10, is designed in a very general way, making it not only a challenge contribution but also part of the ADDO web service brokering system [1597, 1595, 1596]: In order to provide the functionality of the composition algorithms to other software components, it was made accessible as a Web Service shortly after WSC'06. The web service composer is available for any system where semantic service discovery with the Ontology Web Language for Services (OWL-S) [1592] or similar languages is used. Hence, this contest application is indeed also a real-world application.




Fig. 21.10: The WSC 2007 Composition System of Bleul et al. [1604; 1605].

An application accesses the composition system by submitting a service request (illustrated by (b)) through its *Web Service Interface*. It furthermore provides the services descriptions and their semantic annotations. Therefore, WSDL and XSD formatted files as used in the WSC challenge or OWL-S descriptions have to be passed in ((a1) and (a2)). These documents are parsed by a fast *SAX-based Input Parser* (c). The composition process itself is started by the *Strategy Planer* (d). The Strategy Planer chooses an appropriate composition algorithm and instructs it with the composition challenge document (e).

The software modules containing the basic algorithms all have direct access to the *Knowledge Base* and to the *Service Register*. Although every algorithm and composition strategy is unique, they all work on the same data structures. One or more composition algorithm modules solve the composition requests and pass the solution to a *SAX-based Output Writer*, an XML document generating module (f) faster than DOM serialization. Here it is also pos-

sible to transform it to, for example, BPEL4WS [1610] descriptions. The result is afterwards returned through the Web Service Interface (*g*).

One of the most important implementation details is the realization of the operation “promising” since it is used by all composition algorithms in each iteration step. Therefore, we transparently internally merge the knowledge base and the service registry. This step is described here because it is very crucial for the overall system performance.




Fig. 21.11: The Knowledge Base and Service Registry of our Composition System.

A semantic concept is represented by an instance of the class **Concept**. Each instance of **Concept** holds a list of services that directly produce a parameter annotated with it as output. The method `getPromisingServices(A)` of **Concept**, illustrated in Figure 21.11, additionally returns all the **Services** that provide a specialization of the concept *A* as output. In order to determine this set, all the specializations of the concept have to be traversed and their promising services have to be accumulated. The crux of the routine is that this costly traversal is only performed once per concept. Our experiments substantiated that the resource *memory*, even for largest service repositories, is not a bottleneck. Hence, `getPromisingServices` caches its results.

This caching is done in a way that is thread-safe on one hand and does not need any synchronization on the other. Each instance *x* of **Concept** holds an internal variable `promisingServices` which is initially `null`. If `x.getPromisingServices()` is invoked, it first looks up if `x.promisingServices` is `null`. If so, the list of promising services is computed, stored in `x.promisingServices`, and returned. Otherwise, `x.promisingServices` is returned directly. Since we do not synchronize this method, it may be possible that the list is computed concurrently multiple times. Each of these computations will produce the same result. Although all parallel invocations of `x.getPromisingServices()` will return other lists, their content is the same. The result of the computation finishing last will remain in `x.promisingServices` whereas the other lists will get lost and eventually be freed by the garbage collector. Further calls to `x.getPromisingServices()` always will yield the same, lastly stored, result. This way, we can perform caching which is very important for the performance and spare costly synchronization while still granting a maximum degree of parallelization.

Conclusions

In order to solve the 2006 and 2007 Web Service Challenges, we utilized three different approaches, an uninformed search, an informed search, and a genetic algorithm. The uninformed search proofed generally unfeasible for large service repositories. It can only provide a good performance if the resulting compositions are very short.

However, in the domain of web service composition, the maximum number of services in a composition is only limited by the number of services in the repositories and cannot be approximated by any heuristic. Therefore, any heuristic or meta-heuristic search cannot be better than the uninformed search in the case that a request is sent to the composer which cannot be satisfied. This is one reason why the uninformed approach was kept in our system, along with its reliability for short compositions.

Superior performance for all test sets could be obtained by utilizing problem-specific information encapsulated in a fine-tuned heuristic function to guide a greedy search. This approach is more efficient than the other two tested variants by a magnitude.

Genetic algorithms are much slower, but were also always able to provide correct results to all requests. To put it simple, the problem of semantic composition as defined in the context of the WSC is not complicated enough to fully unleash the potential of evolutionary algorithms. They cannot cope with the highly efficient heuristic used in the greedy search. We anticipate however, that, especially in practical applications, additional requirements will be imposed onto a service composition engine. Such requirements could include quality of service (QoS), the question for optimal parallelization, or the generation of complete BPEL [1611] processes. In this case, heuristic search will most probably become insufficient but genetic algorithms and Genetic Programming [65, 1612] will still be able to deliver good results.

Real-World Applications

In this chapter, we will explore real-world applications of global optimization techniques. Such applications are well-researched and established to a point where people are willing to bet money on them. They can safely be utilized in a productive system. Some of these areas where global optimization algorithms are applied in a practical fashion, aiding scientists and engineers with their work, are summarized in this chapter.

22.1 Symbolic Regression

In statistics, regression analysis examines the unknown relation $\varphi : \mathbb{R}^n \mapsto \mathbb{R}$ of a dependent variable $y \in \mathbb{R}$ to specified independent variables $\mathbf{x} \in \mathbb{R}^m$. Since φ is not known, the goal is to find a reasonable good approximation ψ^* .

Definition 22.1 (Regression). Regression¹ [1613, 1614, 1615, 1616] is a statistic technique used to predict the value of a variable which is dependent one or more independent variables.

The result of the regression process is a function $\psi^* : \mathbb{R}^m \mapsto \mathbb{R}$ that relates the m independent variables (subsumed in the vector $\mathbf{x}$ to one dependent variable $y \approx \psi^*(\mathbf{x})$. The function ψ^* is the best estimator chosen from a set Ψ of candidate functions $\psi : \mathbb{R}^m \mapsto \mathbb{R}$. Regression is strongly related to the estimation theory outlined in Section 27.6 on page 450. In most cases, like linear² or nonlinear³ regression, the mathematical model of the candidate functions is not completely free. Instead, we pick a specific one from an array of parametric functions by finding the best values for the parameters.

Definition 22.2 (Symbolic Regression). Symbolic regression [65, 811, 812, 1617, 1618, 1619, 1620, 1621, 813] is one of the most general approaches to regression. It is not limited to determining the optimal values for the set of parameters of a certain array of functions. Instead, regression functions can be constructed by combining elements of a set of mathematical expressions, variables and constants.

22.1.1 Genetic Programming: Genome for Symbolic Regression

One of the most widespread methods to perform symbolic regression is to apply Genetic Programming. Here, the candidate functions are constructed and refined by an evolutionary process. In the following we will discuss the genotypes (which are also the phenotypes) of the evolution as well as the objective functions that drive it. As illustrated in Figure 22.1, the solution candidates, i. e., the candidate functions, are represented by a tree of mathematical expressions where the leaf nodes are either constants or the fields of the independent variable vector $\mathbf{x}$.

¹ http://en.wikipedia.org/wiki/Regression_analysis [accessed 2007-07-03]

² http://en.wikipedia.org/wiki/Linear_regression [accessed 2007-07-03]

³ http://en.wikipedia.org/wiki/Nonlinear_regression [accessed 2007-07-03]




Fig. 22.1: An example genotype of symbolic regression of with $x = \mathbf{x} \in \mathbb{R}^1$.

The set Ψ of functions ψ that can possibly be evolved is limited by the set of expressions Σ available to the evolutionary process.

$$\Sigma = \{+, -, *, /, \exp, \ln, \sin, \cos, \max, \min, \dots\} \quad (22.1)$$

Another aspect that influences the possible results of the symbolic regression is the concept of constants. In general, constants are not really needed since they can be constructed indirectly via expressions. The constant 2.5, for example, equals the expression $\frac{x}{x+x} + \frac{\ln x * x}{\ln x}$. The evolution of such artificial constants, however, takes rather long. Koza [65] has therefore introduced the concept of ephemeral random constants.

Definition 22.3 (Ephemeral Random Constants). If a new individual is created and a leaf in its expression-tree is chosen to be an ephemeral random constant, a random number is drawn uniformly distributed from a reasonable interval. For each new constant leaf, a new constant is created independently. The values of the constant leaves remain unchanged and can be moved around and copied by crossover operations.

According to Koza's idea ephemeral random constants remain unchanged during the evolutionary process. In our work, it has proven to be practicable to extend his approach by providing a mutation operation that changes the value c of a constant leaf of an individual. A good policy for doing so is by replacing the old constant value c_{old} by a new one c_{new} which is a normally distributed random number with the expected value c_{old} (see Definition 27.66 on page 461):

$$c_{new} = \text{random}_n(c_{old}, \sigma^2) \quad (22.2)$$

$$\sigma^2 = e^{-\text{random}_u(0,10)} * |c_{old}| \quad (22.3)$$

Notice that the other reproduction operators for tree genomes have been discussed in detail in Section 4.3 on page 132.

22.1.2 Sample Data, Quality, and Estimation Theory

In the following elaborations, we will reuse some terms that we have applied in our discussion on likelihood in Section 27.6.1 on page 451 in order to find out what measures will make good objective functions for symbolic regression problems.

Again, we are given a finite set of sample data A containing $n = |A|$ pairs of $(\mathbf{x}_i, y_i)$ where the vectors $\mathbf{x}_i \in \mathbb{R}^m$ are known inputs to an unknown function $\varphi : \mathbb{R}^m \mapsto \mathbb{R}$ and the scalars y_i are its observed outputs (possible contaminated with noise and measurement errors subsumed in the term η_i , see Equation 27.237 on page 451). Furthermore, we can access a (possible infinite large) set Ψ of functions $\psi : \mathbb{R}^m \mapsto \mathbb{R} \in \Psi$ which are possible estimators of φ . For the inputs $\mathbf{x}_i$, the results of these functions ψ deviate by the estimation error ε (see Definition 27.52 on page 450) from the y_i .

$$y_i = \varphi(\mathbf{x}_i) + \eta_i \quad \forall i \in [0..n-1] \quad (22.4)$$

$$y_i = \psi(\mathbf{x}_i) + \varepsilon_i(\psi) \quad \forall \psi \in \Psi, i \in [0..n-1] \quad (22.5)$$

In order to guide the evolution of estimators (in other words, for driving the regression process), we need an objective function that furthers solution candidates that represent the sample data A and thus, resemble the function φ , closely. Let us call this “driving force” quality function.

Definition 22.4 (Quality Function). The quality function $f(\psi, A)$ defines the quality of the approximation of a function φ by a function ψ . The smaller the value of the quality function is, the more precisely is the approximation of φ by ψ in the context of the sample data A .

Under the conditions that the measurement errors η_i are uncorrelated and are all normally distributed with an expected value of zero and the same variance (see Equation 27.238, Equation 27.239, and Equation 27.240 on page 451), we have shown in Section 27.6.1 that the best estimators minimize the mean square error MSE (see Equation 27.253 on page 453, Definition 27.59 on page 454 and Definition 27.55 on page 450). Thus, if the source of the values y_i complies at least in a simplified, theoretical manner with these conditions or even is a real measurement process, the square error is the quality function to choose.

$$f_{\sigma \neq 0}(\psi, A) = \sum_{i=0}^{\text{len}(A)-1} (y_i - \psi(\mathbf{x}_i))^2 \quad (22.6)$$

While this is normally true, there is one exception to the rule: The case where the values y_i are no measurements but direct results from φ and $\eta = 0$. A common example for this situation is if we apply symbolic regression in order to discover functional identities [978, 1018, 65] (see also Section 22.1.3). Different from normal regression analysis or estimation, we then know φ exactly and want to find another function ψ^* that is another, equivalent form of φ . Therefore, we will use φ to create sample data set A beforehand, carefully selecting characteristic points $\mathbf{x}_i$. Thus, the noise and the measurement errors η_i all become zero. If we would still regard them as normally distributed, their variance s^2 would be zero, too.

The proof for the statement that minimizing the square errors maximizes the likelihood is based on the transition from Equation 27.248 to Equation 27.249 on page 453 where we cut divisions by s^2 . This is not possible if σ becomes zero. Hence, we may or may not select metrics different from the square error as quality function. Its feature of punishing larger deviation stronger than small ones, however, is attractive even if the measurement errors become zero. Another metric which can be used as quality function in these circumstances are the sums of the absolute values of the estimation errors:

$$f_{\sigma=0}(\psi, A) = \sum_{i=0}^{\text{len}(A)-1} |y_i - \psi(\mathbf{x}_i)| \quad (22.7)$$

22.1.3 An Example and the Phenomenon of Overfitting

If multi-objective optimization can be applied, the quality function should be complemented by an objective function that puts pressure in the direction of smaller estimations ψ . In symbolic regression by Genetic Programming, the problem of code bloat (discussed in Section 4.10.4 on page 188) is eminent. Here, functions do not only grow large because they include useless expressions (like $\frac{x*x+x}{x} - x - 1$). A large function may consist of functional expressions only, but instead of really representing or approximating φ , it is degenerated to

just some sort of misfit decision table. This phenomenon is called overfitting and has initially been discussed in Section 1.4.6 on page 48.

Let us, for example, assume we want to find a function similar to Equation 22.8. Of course, we would hope to find something like Equation 22.9.

$$y = \varphi(x) = x^2 + 2x + 1 \quad (22.8)$$

$$y = \psi^*(x) = (x + 1)^2 = (x + 1)(x + 1) \quad (22.9)$$

i	x_i	$y_i = \varphi(x_i)$	$f_2^*(x_i)$
01	-5	16	15.59
1	-4.9	15.21	15.40
2	0.1	1.21	1.11
3	2.9	15.21	15.61
4	3	16	16
5	3.1	16.81	16.48
6	4.9	34.81	34.54
7	5	36	36.02
8	5.1	37.21	37.56

Table 22.1: Sample Data $A = \{(x_i, y_i) : i \in [0..8]\}$ for Equation 22.8

For testing purposes, we choose randomly the nine sample data points listed in Table 22.1. As result of Genetic Programming based symbolic regression we may obtain something like Equation 22.10, outlined in Figure 22.2, which represents the data points quite precisely but has nothing to do with the original form of our equation.

$$\begin{aligned} \psi_2^*(x) = & ((((((0.934911896352446 * 0.258746335682841) - (x * ((x / ((x - \\ & 0.763517999368926) + (0.0452368900127981 - 0.947318140392111)))) / ((x - (x + x)) + \\ & (0.331546588012695 * (x + x)))))) + 0.763517999368926) + ((x - (((0.934911896352446 * \\ & ((0.934911896352446 / x) / (x + 0.947390132934724)))) + (((x * 0.235903629190878) * (x - \\ & 0.331546588012695)) + ((x * x) + x))) / x) * (((x - (x * (0.258746335682841 / \\ & 0.455160839551232))) / (0.0452368900127981 - 0.763517999368926)) * x) * \\ & (0.763517999368926 * 0.947318140392111)))) - (((((x - (x * (0.258746335682841 / \\ & 0.455160839551232))) / (0.0452368900127981 - 0.763517999368926)) * 0.763517999368926) \\ & * x) + (x - (x * (0.258746335682841 * 0.934911896352446)))))) \end{aligned} \quad (22.10)$$

We obtained both functions ψ_1^* (in its second form) and ψ_2^* using the symbolic regression applet of Hannes Planatscher which can be found at <http://www.potschi.de/sr/> [accessed 2007-07-03]⁴. It needs to be said that the first (wanted) result occurred way more often than absurd variations like ψ_2^* . But indeed, there are some factors which further the evolution of such eyesores:

- If only few sample data points are provided, the set of prospective functions that have a low estimation error becomes larger. Therefore, chances are that symbolic regression provides results that only match those points but differ in all other points significantly from φ .
- If the sample data points are not chosen wisely, their expressiveness is low. We for instance chose 4.9, 5, and 5.1 as well as 2.9, 3 and 3.1 which form two groups with members very close to each other. Therefore, a curve that approximately hits these two clouds is rated automatically with a high quality value.

⁴ Another good applet for symbolic regression can be found at <http://alphard.ethz.ch/gerber/approx/default.html> [accessed 2007-07-03]




Fig. 22.2: $\varphi(x)$, the evolved $\psi_1^*(x) \equiv \varphi(x)$, and $\psi_2^*(x)$.

- A small population size decreases the diversity and furthers “incest” between similar solution candidates. Due to a lower rate of exploration, only a local minimum of the quality value will often be yielded.
- Allowing functions of large depth and putting low pressure against bloat (see Section 4.10.4 on page 188) leads to uncontrolled function growth. The real laws φ that we want to approximate with symbolic regression do usually not consist of more than 40 expressions. This is valid for most physical, mathematical, or financial equations. Therefore, the evolution of large functions is counterproductive in those cases.

Although we made some of these mistakes intentionally, there are many situations where it is hard to determine good parameter sets and restrictions for the evolution and they occur accidentally.

22.1.4 Limits of Symbolic Regression

Often, we cannot obtain an optimal approximation of φ , especially if φ cannot be represented by the basic expressions available to the regression process. One of these situations has already been discussed before: the case where φ has no closed arithmetical expression. Another possibility is that the regression method tries to generate a polynomial that approximates the φ , but φ does contain different expressions like $\sin$ or e^x or polynomials of an order higher than available. Yet another problem is that the values y_i are often not results computed by φ directly but could, for example, be measurements taken from some physical entity and we want to use regression to determine the interrelations between this entity and some known parameters. Then, the measurements will be biased by noise and systematic measurement errors. In this situation, $f(\psi^*, A)$ will be greater than zero even after a successful regression.

Research Applications

Research applications differ from real-world application by the fact that they have not yet reached the maturity to be applied in the mainstream of their respective area. They initiate a process of improvement and refinement, until we obtain solutions that are on par or at least comparable with those obtained by the traditional methodologies. Such a process can, for instance, be observed when following the progress in the area of Genetic Programming via the book series of Koza [65, 925, 926, 850]. On the other hand, research applications differ from toy problems because they are not intended to be used as sole demonstration examples or benchmarks but are first steps into a new field of application.

The future of a research application is either to succeed and become a real-world application or to fail. In case of a failure, it may turn into a toy application where some certain features of global optimization methods like evolutionary algorithms can be tested.

Currently, I am not so happy with the division in *research* and *real world* applications (see previous chapter) because many of the approaches listed in this chapter are already very sophisticated. Nevertheless, I would like to somehow distinguish between projects still developed for scientific purposes and those which already became part of industrial and economic procedures. Then again, I also group applications together that belong to one topic, like I do in Section 23.1, for instance. And, of course, I cannot always know to which extend one approach has already been utilized. Thus, you may find some projects listed in this chapter that maybe would fit better into the next one or vice versa. Anyway, with this classification I do not want to belittle the work of fellow researchers but instead separate promising new ideas from already widely established ones.

23.1 Global Optimization of Distributed Systems

In this and the following chapter we will discuss different applications of global optimization for improving the features of distributed systems. In the context of this book, we divide these applications into two groups: On one hand, there are Global optimization approaches for optimizing the network structure and management or the message flow. On the other hand, we can try to utilize Genetic Programming for translating a desired global behavior of a network into the algorithms that need to be executed locally on its nodes in order to achieve that behavior [861, 862].

One of the differences between the two groups is the hardness of the problems they impose. Genetic Programming of distributed algorithms suffers from highly rugged objective functions, epistatic genomes¹, and weak causality. Such problems occur less intense in the area of global optimization of network structures. Whereas these approaches have already been successfully applied in research and practice, we still need to research whether Genetic

¹ See Section 4.8 on page 169 for a discussion about epistasis in Genetic Programming.

Programming of distributed algorithms is feasible and whether or not it can provide us with efficient and robust solutions.

In the following, we will first outline some of the research done in global optimization of network components and structures and later focus on Genetic Programming of distributed algorithms.

Optimization algorithms are methods for finding optimal configurations. In distributed systems, there are lots of things that can be configured somehow, like the topology, security, and routing. Hence, there is a giant potential for using global optimization in order to improve networks. And this potential is widely utilized. The study of Sinclair [480] from 1999 reported more than 100 papers had been published that utilized Evolutionary Computation for optimizing network topology and dimension, node placement, trees, routing, and wavelength and frequency allocation. These topics are also discussed in the book of Corne et al. [277] from 2000. Now these works are already about ten years old but recent studies like the one of Alba and Chicano [486] and Cortés Achedad et al. [481] as well as many recently published papers show that the interest in applying global optimization techniques in this research field has by no means decreased.

23.1.1 Optimizing Network Topology and Dimensioning

The topology of a network is a graph that defines which of the nodes are able to communicate with each other.² The structure of any network, be it a distributed system, a social network, or a transportation system, can be described as such a graph. Especially for communication networks, the topology is very important since it defines the maximum throughput, the latency and hops between the nodes, the robustness in terms of link failures, the installation costs, etc.. And these are also exactly the objectives which can be optimized when global optimization algorithms are used to design the network topology.

The application of Evolutionary Computation in this area has a very long tradition, starting with the work of Michalewicz [1622] and Kumar et al. [714] in 1991 and 1992. Before the year 2000, already more than 30 papers had been published in this area [872] and their number is still increasing. In this section, we will outline some of these methods for finding good network structures. Some of the other researchers that have contributed optimization approaches for network topologies are Pierre [1623], [715, 1383, 716], Brittain et al. [718], Ko et al. [719], Dengiz et al. [720], and Coombs and Davis [1624]. The closely related topic of assigning cells to transmitters/switches in wireless transmission has also been extensively studied, for instance, by Pierre and Houéto [1385; 1384], Quintero and Pierre [717; 1448; 484], and Alba et al. [485]. More related work can be found in Sinclair [480].

Solving the Terminal Assignment Problem

The terminal assignment (TA) problem involves determining the minimum cost links to connect a given set of nodes (the terminals) to another (disjoint) set of nodes (the concentrators). The required capacity of each terminal is known and may vary from terminal to terminal. The capacity of the concentrators is also known and so are the costs for linking them to the terminals. Each terminal must be linked to exactly one concentrator in a way such that the maximum capacity of no concentrators is exceeded. An assignment has to be found under the objective of minimizing the total costs. This problem is $\mathcal{NP}$ -hard [1625], except for the special case where all terminals have the same capacity requirements and all concentrators have the same capacity.

Abuali et al. [1626] first applied genetic algorithms to instances of this problem where the capacity of all concentrators was equal. Khuri and Chiu [709], investigated the utility of different heuristic algorithms, including greedy approaches and genetic algorithms for the general case where the capacities of the single concentrators differ. In their first genetic

² See Section 29.2.1 on page 486 for more information on network topologies.

approach, Khuri and Chiu [709] use a terminal-centric representation in form of a string genome $G = X = \mathbb{N}^n$ consisting of natural numbers, where n is the number of terminals. The value of an element s_i of a genotype/phenotype string $S = (s_1, s_2, \dots, s_n)$ stands for the concentrator to which the i^{th} terminal is connected. Strings that were infeasible, i. e., violate at least one of the constraints given above, are penalized with an offset separating them from all feasible solutions. Additionally, they also tested a genetic algorithm with a genome especially designed for grouping of elements. The results of their work indicated that the first genetic algorithm provided better results than the greedy approach which, in turn, outperformed the grouping GA.

Yao et al. [1447] apply a hybrid evolutionary algorithm to a multi-objective variant of this problem where, besides the total cost, also the number of concentrators used is to be minimized. They also abandon the terminal-centric representation in favor of a concentrator-centric genome consisting of m trees of depth 1, where m is the number of concentrators. Each concentrator is represented as root of a tree and terminals are leaves linked to these roots. If a root has no children, then the corresponding concentrator is not used. The experimental results of Yao et al. [1447] showed that evolutionary algorithms hybridized with Lamarckian or Baldwin local search both perform approximately equally well and are suitable for solving the terminal assignment problem. The concentrator-based representation also showed to be advantageous in these experiments compared with the terminal-centric genome.

Singular, Selfish, and Self-Organized (S^3) Networks

Zweig [1627] defines an iterative online approach for self-organizing of networks of selfish nodes. Her algorithm is not population-based – at point in time there is exactly one (singular) network. In each time step, one node may modify the edges to its neighbors by either adding one connection (from itself) to another node or removing one edge attached to it.

If this modification seems to be advantageous from the perspective of the node (therefore *selfish*), it is committed, otherwise it is rejected. As measure for the utility of a modification, a node may, for instance, compute the maximum distance from itself to any other node or the sum of the distances to all other nodes. In [454], it is shown that self-organization algorithms minimizing such objectives can decrease the diameter of a network down to a certain threshold in polynomial time.

Zweig [1627] refers to this approach evolutionary algorithm. Since its population has the size one and only a unary reproduction operator is used, one may rather consider it as a hill climbing approach (which you can find discussed in Chapter 8 on page 209).

Optimization in the European Optimal Network

With the COST Action 239 [1628], the European Commission studied the feasibility of the European Optical Network (EON), a transparent optical overlay network capable of carrying all the international traffic between twenty of the main centers of Europe [871]. The initial design for this network is illustrated in Figure 23.1.

Sinclair [712] evolved a network topology for EON using a genetic algorithm with an objective function that took the estimated traffic on the links between the nodes and weights of the nodes themselves into account. This way, he was able to achieve significant improvements of both, costs and reliability, compared to the topology initially design [1629]. The same group of researchers later applied a similar GA to also optimize the route and wave length assignment for such a network. Again, they obtained results superior to initial expectations [713].

In [871] and [872], they tried out four different Genetic Programming approaches for the topology optimization problem and compared them with the initial genetic algorithm. We will outline the representations used in these approaches on the example (also used in [712]) of the $N = 9$ central nodes (Amsterdam, Berlin, Brussels, Luxembourg, London, Milano,




Fig. 23.1: The initial network design of the EON network [712, 1629].

Paris, Prague, and Zurich). In Figure 23.1, the black dots stand for these central nodes and the black lines for the connections between them.

The Genetic Algorithm

In the genetic algorithm, Sinclair used a genome where each of the $L = 0.5N(N - 1) = 0.5 * 9 * 8 = 36$ possible links between these nine metropolises was represented by one bit. The central topology from Figure 23.1 thus translates to the genotype in Figure 23.2 (which effectively encodes the connection matrix).




Fig. 23.2: The genotype representing the connection matrix of the nine central nodes of EON.

Genetic Programming with a Relational Function Set

With the relational function set (RFS), a function `isLinkThere(LID)` was evolved which is 1 if the link LID is present and 0 otherwise. Therefore, each of the $L = 36$ possible links is associated with a unique number from $[0..L - 1]$. The terminal symbols in this approach are the link ID LID (i. e., the parameter of `isLinkThere`) and constants. As function set, addition

(+) and subtraction (−) *modulo* L are provided together with greater than (>) and less than (<) which return 1 if their first argument greater or lesser than their second one. The decision function *if* returns the value of its second argument if the first one equals to 1 and the value of the third one otherwise. Since the function evolved could return values from 0 to $L - 1$, it was wrapped in a way that mapped 0.5 to 1 and the rest to 0. Fig. 23.3.1 sketches how such a symbolic regression-like GP representation of the bit string from Figure 23.2 may look like, using the LIDs as defined there.




Fig. 23.3.1: The RFS approach.




Fig. 23.3.2: The decision tree method (DT).




Fig. 23.3.3: The connective approach (CN).




Fig. 23.3.4: The node-pair encoding GP (NP2).

Fig. 23.3: Snippets from the GP representations of the bit string from Figure 23.2

Genetic Programming for Decision Trees

In their DT-Genetic Programming approach, Aiyarak, Saket, and Sinclair [871] extended their previous method by combining the < and > with the *if* statement. With this modification, they wanted to evolve decision trees and hoped to reduce the computational effort for finding solutions. For illustration purposes, we have translated the program from Fig. 23.3.1 to this representation and depict the result in Fig. 23.3.2.

Genetic Programming – the Connective Approach

The third Genetic Programming approach (called *connected nodes*, or CN for short) differs very much from the other two. As terminals, it only uses node identifiers and it only has one function, *con* [872]. *con* simply adds a link between the two nodes identified by its parameters and returns its first argument. The tree in Fig. 23.3.3 describes how the network defined in Figure 23.2 can be built.

Node-Pair encoding GP

Inspired by the *edge encoding* method by Luke and Spector (discussed in Section 4.4.4 on page 144), Sinclair went on and designed another Genetic Programming approach to the problem. His *node-pair encoding* (NP) involves a tree-to-graph genotype-phenotype mapping which explicitly constructs the network [872, 873]. Table 23.1 illustrates the function and terminal sets of both variants of NP which only differ in the number of children of the function nodes (noted in the columns NP1 and NP2). NP programs work with tuples of node IDs (a, b) with $a, b \in [0..N - 1]$ (and $N = 9$ for our example). The program tree is processed from top to bottom and begins with the input $(0, 0)$ into the root node. Each node can either manipulate the value of the first element of the tuple (with **da**, **ia**, **ia2**, ...), rotate the tuple (**rev**), create or cut a connection between the two nodes in the tuple (**add** or **cut**), or simply does nothing (**nop**). If it has children, it will pass the (resulting) node tuple to all its children. Fig. 23.3.4 illustrates how an NP2-version program snippet is executed and the values of the node tuples at the different levels.

Abbreviation	NP1	NP2	Description
rev	1	1	reverse: $(a, b) \rightarrow (b, a)$
da	1	2	modulo-decrement a : $(a, b) \rightarrow (a - 1 \bmod N, a)$
ia	1	2	modulo-increment a : $(a, b) \rightarrow (a + 1 \bmod N, a)$
ia2	1	2	modulo-increment a by 2: $(a, b) \rightarrow (a + 2 \bmod N, a)$
ia4	1	2	modulo-increment a by 4: $(a, b) \rightarrow (a + 4 \bmod N, a)$
ia8	1	2	modulo-increment a by 8: $(a, b) \rightarrow (a + 8 \bmod N, a)$
dbl	2	2	plain doubling: pass node pair to both children
tpl	3	3	plain tripling: pass node pair to all three children
add	1	0	if $a \neq b$ create link between a and b
cut	1	0	if a and b are linked, remove the link
nop	0	0	do nothing (terminal)

Table 23.1: The Node-Types of the NP-approaches.

Summary

In [871], Aiyarak et al. tested the first three GP approaches (RFS, DT, and CN) and found that none of them produced solutions with better fitness than the original genetic algorithm method. RFS-GP performed worst. The DT method was better than CN on the small test set containing only the nine central nodes and scored even with the GA. On the complete EON network, CN in turn outperformed DT but was still worse than the solution provided by the genetic algorithm. All Genetic Programming approaches needed much (at least five times) more individual evaluations to find their solution than the genetic algorithm. These results were confirmed in [872], where CN and the two versions of NP were compared with the GA. Again, the GA was the best approach, but NP2 found the best result for one of the five test problems and came close to the GA in all other runs. Still, the experiments with Genetic Programming needed much longer than those with the genetic algorithm. The lesson from these experiments is that representation is very important in optimization. Sometimes a simple, fixed-length bit string can outperform more complex approaches. And even if one class of algorithms (like GP) is applied, modifications in the representation used can still lead to vast differences in performance.

23.1.2 Optimizing Routing

Routing is the way in which messages are relayed over multiple nodes in a network in order to reach their destination. Naturally, there often exist multiple different routes from a source

to a destination. The objective of a good routing algorithm is to guide the messages along the shortest possible routes. A path in the network gets blocked if any of its nodes is congested, i. e., has to process too many message at once. Additional paths may be added to the net as new nodes are connected, on the other hand. Hence, the routing algorithm has to avoid congestion and should incorporate new links as soon as possible in order to increase the network utility.

In this section, we will discuss some interesting approaches for improving routing with global optimization algorithms. Other contributions to this area have been provided by, for instance, Louis Anthony Cox, Jr. and Qiu [1630], Xiang et al. [721], Wang et al. [722], Zhu et al. [723] and Galiasso and Wainwright [724].

Improving Routing in Gnutella

The Gnutella network [1631, 1632] is one of the oldest and simplest peer-to-peer³ file sharing systems. In Gnutella, each node only interacts with neighboring nodes which it chooses at startup. Normally, a node picks seven such neighbors and tries to maintain connections to them. Iles and Deugo [868] use Genetic Programming for improving routing in such networks. They evolve rules for dynamically setting the number of neighbor nodes and for selecting (and possibly exchanging) them. These rules use information the peer-to-peer node can aggregate over time, like the (current) number of neighbors, the bandwidth of the node, the bandwidth of the neighbor nodes, and the duration of the connections to them. The expressions produced by Genetic Programming led to significant performance improvements in simulations. The results of the experiments of Iles and Deugo [868] still use a fixed number of neighbors, but tend to prefer five or six partner nodes instead of Gnutella's default of seven. The way these neighbors are picked, however, becomes a dynamic process depending on multiple aggregated parameters such as the distance in hops and number of search hits to the candidate nodes.

Evolving Fault-tolerant Routing Rules

The problem of finding the shortest paths from one vertex in a graph to all others has been solved in 1959 by Dijkstra⁴ [1633]. Forty or so years later, Kirkwood et al. [869], [870] use Genetic Programming to obtain such paths by evolution. This approach is not in competition with Dijkstra's algorithm, since it aims to breed robust routing rules for networks where links may fail.

The input of the system of Kirkwood et al. [869] is the ID of the target node and a graph (containing this node) where each node has a unique ID and each connection has a certain length. Genetic Programming breeds S-expression which are nested instances of the domain-specific **IF-CUR-GO** instruction. The four parameters **W X Y Z** of this instruction each stand for a node name and its return value is again, a node name. On each node a message reaches when travelling through the network described by the graph, the routing rule is evaluated and it is sent to the node whose ID is returned. **IF-CUR-GO W X Y Z** evaluates to **X** if the current node is **W**, **X** and **W** are directly connected, and **X** has not yet been visited twice. It returns **Y** if one of the latter two conditions is violated and **Z** if the current node is not **W**.

As objective function, Kirkwood et al. [869] set the difference in length between the optimal route (computed with Dijkstra's algorithm) and the evolved route. As result, they obtain rules that are very robust in terms of link failures that, for instance, find correct routes in two third of all cases even if 47% of the connections have broken down. In Figure 23.4, a routing rule set resulting from one of the experiments of Kirkwood et al. [869] is displayed on the left. It routes messages through the graph sketched on the right-hand side to node 0.

³ See Definition 29.30 on page 489 for a discussion of peer-to-peer networks.

⁴ http://en.wikipedia.org/wiki/Dijkstra's_algorithm [accessed 2008-06-24]




Fig. 23.4: An evolved routing rule set (left) with target node 0 for the specified network (right).

Genetic Adaptive Routing Algorithm

In the internet, messages are forwarded to their destinations on the basis of routing tables. These tables can be filled with, for example, Dijkstra's algorithm. They contain either the address of the next hop, i. e., the next computer to send a packet to, or the complete routes to each possible destination for source routing. Protocols like RIP [1634] and OSPF [1635] build and distribute such tables, but disregard dynamic aspects like communication latency.

Munetomo et al. [1636; 1637; 1638], [1639] contributed GARA, the genetic adaptive routing algorithm, for building and maintaining routing tables for source routing. GARA runs independently on every node in the network. Each node maintains a list of multiple complete routes to the most frequently contacted communication partners. After a specific number of packets have been sent along a route, an additional packet is sent for observing the communication latency. The fitness of the routes and their likeliness for being chosen then corresponds to this delay. In Figure 23.5, the table with the destinations most frequently addressed by node 0 is sketched in the context of an example network.




Fig. 23.5: The table with the most frequently used routes of node 0.

From time to time, *path genetic operators* such as path mutation and path crossover, both illustrated in Figure 23.6, are applied to routes with the same destination. With these operators, new paths are discovered and made available for message propagation. Path mutation therefore simply exchanges a node within a route with another one (Fig. 23.6.1). Path crossover (Fig. 23.6.2) is applied to two routes that pass through one common point (other than their source and destination). The sub-paths before and after that point then are exchanged. Selection is performed when the routing table exceeds a specific size. By utilizing these evolutionary primitives, GARA becomes a self-organizing system able to




Fig. 23.6.1: The path mutation operator.




Fig. 23.6.2: The path mutation operator.

Fig. 23.6: The path genetic operations in GARA.

adapt to changing network situations and to discover new routes challenging established protocols on equal footing.

Ant-based Routing

The routing of network messages or phone calls through a telecommunication system can be compared with the transportation of food by ants to their nest [1410].

Schoonderwoerd et al.

Schoonderwoerd et al. [1397; 1640], [1409] have developed an routing method based on this paragon in form of an online Ant Colony Optimization⁵ approach. In their systems, the nodes exchange messages (the ants) from time to time in order to build and maintain “pheromone tables” used for routing. In principle, a pheromone table of a node with m neighbors in a network N with $n = \text{len}(N)$ nodes is a $n \times m$ matrix. For any of the n possible destinations, it contains a probability that a message is routed to any of the m neighbors. At each time step during Schoonderwoerd’s simulations, any node could launch an ant with a random

⁵ See Chapter 13 on page 233 for more information on Ant Colony Optimization.

destination. Whenever an ant created by node i arrives at node j coming from node k , it will update the pheromone table P_j of j according to

$$P_{j[i,k]} = \frac{P_{j[i,k]} + \Delta}{1 + \Delta} \quad (23.1)$$

$$\forall l \in [1..n], l \neq k \quad P_{j[i,l]} = \frac{P_{j[i,l]}}{1 + \Delta} \quad (23.2)$$

In other words, the probability of selecting node k as next hop for future messages with destination i is increased while decreasing the probabilities of the other possible next hops of such a message. The total probability of all hops together stays 1. This update method works in the opposite direction of the ants movements. It indirectly influences the forward probability of the path taken by the ant since subsequent ants moving to node i along it will again have backward influence which then points into the original forward direction.

Schoonderwoerd adds an “age” to the ants which is increased in every hop in order to foster the usage of short pathways. Δ becomes a function which decreases with the age of the ant. Thus, the older an ant, i. e., the longer the path it has travelled, the lower will its influence of guiding future ants along this path be. Congested relay stations can artificially increase the age of the ants passing them and may also delay the ants on their way. This will decrease their traffic.

New routes can be discovered by introducing a factor $f : 0 < f < 1$ which denotes a probability that the next hop of an ant is chosen totally randomly whereas the pheromone table is used for the decision with probability $1 - f$.

In summary, Schoonderwoerd et al. [1397] have presented a decentralized and self-organized online optimization process for communication networks which can be generalized to arbitrary routing problems. The simple behavior of their ants leads to interesting, emergent overall configurations of the network it is applied to.

AntNet

Di Caro and Dorigo [1417; 1416], [1414] develop this concept further in their AntNet algorithm. For each node j , it builds a routing table P_j which has basically the same structure as the ones used by Schoonderwoerd et al.⁶ What changed is the way that this matrix is updated.

Regularly, every node s sends an agent (*forward ant*) $F_{s \rightarrow d}$ to a node $d \neq s$ (randomly selected according to the local traffic) in order to discover feasible and efficient paths. Forward ants share the same queues as data packets, so they will experience the same delays [1417]. These ant agents keep a history of their route and of the times needed for its sections.

At each node j , the travelling ant $F_{s \rightarrow d}$ selects the next node to visit according to:

$$P(j \rightarrow i) = \frac{P_{j[d,i]} + \alpha l(i)}{1 + \alpha(m-1)} \quad (23.3)$$

Where $P(j \rightarrow i)$ is the probability that i is chosen as next node and $P_{j[d,i]}$ ⁷ is the routing table entry which describes the expected utility of i as next hop for messages going to d . m is the number of neighbors of node j and $l(i)$ is a value proportional to the queue length for destination i normalized to the unit interval. The influence of $l(i)$ makes the system more reactive to the current traffic situation. It is weighted by the factor α which usually takes on values between 0.2 and 0.5.

⁶ Well, except for naming changes and that the rows and columns are switched. We will ignore these minor differences in this work.

⁷ Again, notice that we have switched the columns and rows in Di Caro and Dorigo’s approach in order to relate it to the one of Schoonderwoerd et al..

If the forward ant $F_{s \rightarrow d}$ detects that it somehow travelled in a loop, it dies. Otherwise, when it arrives at its destination d , it is converted into a backward agent $B_{d \rightarrow s}$ which travels back to s along the same route that also $F_{s \rightarrow d}$ took. When arriving at a node j coming from a node i , the entry for choosing i as next hop for messages going to d in the routing table of j is updated as follows:

$$P_{j[i,d]} = P_{j[i,d]} + r(1 - P_{j[i,d]}) \quad (23.4)$$

where $r \in [0, 1]$ is a reinforcement factor, a measure of goodness of the observed trip time, based on a statistical model for the local traffic M (also updated by the agents) and the time. By multiplying r with $(1 - P_{j[i,d]})$, small probabilities are increased proportionally more, which allows faster exploration of new, good routes. After this update, corresponding column of P_j is normalized again.

Global Information?

The simulations of Di Caro and Dorigo [1417] show that AntNet is an efficient and robust routing algorithm. However, it makes one the assumption that global information is available. Every node holds a routing table with entries defining how to send messages to any other node. Such global information is normally not available and furthermore would also be subject to constant changes when routers are added to or removed from the network.

Liang et al. [1413; 1415] point out this weakness and test a LocalAnt algorithm restricted to routing tables restricted to two entries for each neighboring node (for the neighbor itself and the probability to route over this node for other destinations). They show that, under these more realistic conditions, the efficiency of the algorithm is much lower than in the case where global information is available (this original case is called GlobalAnt in this context). Liang et al. [1415] propose a distributed algorithm where each node holds a population of individuals encoding routes. These individuals travel over the network as genetic agents and, subsequently, as backward agents similar to the AntNet approach and are subject to recombination, mutation, and aging. According to their study [1415], the performance of this method falls in between those of LocalAnt and GlobalAnt.

Genetic Programming of Broadcasting Algorithms

Broadcasting in a graph is the process of spreading information which is initially known to only one node to all other nodes, subject to the following rules:

- Only a node which already knows the information can spread it further.
- A node can only send one message per time step.
- A node can only send a message over vertexes which are connected to it.

Different broadcasting schemes can be defined for a topology, and we are (of course) interested in finding the one which disseminates the information to all nodes in the shortest possible time.

Comellas and Giménez [858] have used Genetic Programming for growing such broadcasting schemes for directed 2-grids, toroidal grids, hypercubes, cube-connected-cycles, and butterfly graphs. For the butterfly graph, a solution even better than the best known upper bound at that time (by Klasing et al. [1641]) was found, whereas in all other cases, (known) optimal algorithms evolved.

Example: Directed 2-Grid

In their experiments, Comellas and Giménez [858] used Standard Genetic Programming for evolving Lisp S-expressions⁸. The function set for the directed 2-grid example was contained the following expressions:

⁸ See Section 29.3.11 on page 512 for more information about S-expressions.

```


1 (Prog3 (If0ri      (Prog2 (IfTurn0 MoveHor)
2                               (IfTurn1 MoveVer)))
3       (IfProcHor (Prog3 (IfTurn0 MoveHor)
4                               (IfTurn0 MoveVer)
5                               (IfTurn1 MoveHor))))
6       (IfProcVer (IfTurn0 MoveVer)))

```

Listing 23.1: The optimal broadcasting scheme evolved for 2-dimensional, directed grids.

- If0ri executes its parameter action if the node the program runs on was the one who initially knew the information to be disseminated.
- IfProcHor and IfProcVer execute their parameter if they received the information from a horizontally or vertically adjacent node.
- Proc1, Proc2, Proc3, and Proc4 concatenate 1, 2, 3, or 4 instructions to a sequential block.
- IfTurn0/IfTurn1 execute their actions only in the first/second time step after the node received the information.

As terminals, `MoveHor` and `MoveVer` were available, which send the information to the next horizontally or vertically adjacent node as well as `NULL` which does nothing. For computing the objective values, the algorithms evolved where applied to a graph. The number of nodes reached was counted, multiplied with a factor T . From this value, 2 are subtracted for each condition value and `NULL` action and 4 for each other action. The connective `ProgN` functions had no influence on fitness. For the 2-dimensional directed grid, Listing 23.1 was yielded as result.

Fig. 23.7: The application of Listing 23.1 to a 5×5 grid.

In Figure 23.7, we have applied this algorithm to a 5×5 grid. The gray nodes have no knowledge about the information and turn black ones they receive it. It is interesting to know that the `IfTurn0 MoveVer` in line 4 of Listing 23.1, although seemingly useless, is important. Without it, the last node in the second line of Figure 23.7 would not be reached in time step 5 but in step 6, since the `MoveHor` action cannot be executed by the last node in the first line. With the `MoveVer` with the same preconditions, an unnecessary pause is avoided.

For the other topologies mentioned, Comellas and Giménez [858] have used modified function and terminal sets and obtained similarly optimal results.

Optimizing Collective Communication

Recently, Jaros and Dvorak [1642] developed a scheduling technique collective communication (CC) on basis of evolutionary algorithms. Their work addresses high-performance multi-processor systems where the single CPUs are connected via NoCs (networks on chip). Jaros and Dvorak experiment with one-to-all broadcast (OAB), all-to-all broadcast (AAB), one-to-all scatter (OAS, a private message to each partner), and all-to-all scatter (AAS) [1642].

Each collective communication process is considered as a set of point-to-point communications and the CC scheduling problem is defined as partitioning this set into as few as possible subsets. The communications in each subset can be executed in parallel and the subsets themselves are executed one after the other in a sequence of synchronized steps. The objective is to find a minimal partition without causing conflicts.

Jaros and Dvorak [1642] introduce different encodings for broadcast and scatter communication as well as special reproduction and repairing operators. The objective function applied corresponds to the number of conflicts that a schedule introduces, i. e., the number of times that two point-to-point communications share the same link in the same time step. The experimental results showed that this approach is very efficient and even led to an improvement of the theoretical lower bounds of the number of steps needed for AAS communication.

23.1.3 Synthesizing Protocols

Protocols like the IP [1643] and TCP [1644] are the rules for message and information exchange in a distributed system. Depending on the application, they can become arbitrarily complex and strongly influence the efficiency and robustness of a distributed system.

Transformation of Petri Net-based Specifications

In [1645], Yamaguchi et al. define the problem of transforming a service specification in form of a Petri Net with registers (PNR) to a protocol specification in the same format. Later, El-Fakihiy et al. [706] show how to solve this problem efficiently with genetic algorithms under the objective of minimizing communication costs.

Transformation of Finite State Machine-based Specifications

A similar approach for synthesizing protocol specifications from service specifications has been contributed by de Araújo et al. [859; 1646; 1647]. We will use the communication protocol example given in [859] to summarize this method in which all specifications are represented as finite state machines.

In Figure 23.8, we sketch the starting situation as given by de Araújo et al. [859] in their example. At the left side, a global service specification is defined in form of a finite state machine which describes the interactions that take place in a communication protocol between a sender and a receiver. It contains only the interactions *primitives* between a user (or a user application) and the peer entities, where *?xyz* represents input from and *!xyz* stands output to the user. Events on the sender side are marked with **S** and those occurring at the receiver are annotated with **R**.

The finite state machine starts in state 1. On the sender side, the user may then issue a connection request **CReq** causing a transition to state 2. A window indicating an incoming connection **CInd** may then pop up at the receiver side which then can either be confirmed ($3 \rightarrow 4$) or declined ($3 \rightarrow 7$). In the later case, a disconnection indication will ensue on the sender side whereas the former transition leads to a connection confirmation. A connection may be reconnected (states 5, 8, 9, 10) or being closed by either the sender ($5 \rightarrow 6$) or




Fig. 23.8: The scenario for protocol synthesis as used in [859].

the receiver (5 → 7) with a disconnect request **DReq**, causing a disconnect indication **DInd** (maybe again in form of a window) to pop up on the other end.

The goal is now to evolve a finite state machine that governs the behavior of one of the two protocol partners and to synthesize the messages PDUs⁹ that have to be exchanged between the sender and the receiver.

de Araújo et al. [859] therefore first create a sequence of interaction primitives by traversing the service specification FSM. In this traversal, all state transitions need to be visited. One such sequence could be (**?S:CReq**, **!R:CInd**, **?R:CResp**, **!S:CConf**, **?S:RReq**, **!R:RInd**, **?R:RResp**, **!S:RConf**, **!S:DReq**, **!R:DInd**, **?S:CReq**, **!R:CInd**, **?R:DReq**, **!S:DInd**, **?S:CReq**, **!R:CInd**, **?R:CResp**, **!S:CConf**, **?R:DReq**, **!S:DInd**). It is obvious that if one event on the sender side is followed by an event on the receiver side (or vice versa) and the second event is not triggered by the user (**?xyz**), a message (i.e., a PDU) has to be exchanged between the two sides. These PDUs can either be automatically generated or predefined and are then inserted into the sequence, resulting in the list, (**?S:CReq**, **pduConReq**, **!R:CInd**, **?R:CResp**, **pduConAcc**, **!S:CConf**, **?S:RReq**, **pduReiReq**, **!R:RInd**, **?R:RResp**, **pduReiAcc**, **!S:RConf**, **!S:DReq**, **pduDisInd**, **!R:DInd**, **?S:CReq**, **pduConReq**, **!R:CInd**, **?R:DReq**, **pduDisInd**, **!S:DInd**, **?S:CReq**, **pduConReq**, **!R:CInd**, **?R:CResp**, **pduConAcc**, **!S:CConf**, **?R:DReq**, **pduDisInd**, **!S:DInd**), for instance.

From this list, the message exchanges concerning the protocol peer which should be evolved are extracted. In [859], that was the sender side, which would here result in the following tuples: ((**?S:CReq**, **pduConAcc**), (**pduConAcc**, **!S:CConf**), (**?S:RReq**, **pduReiReq**), (**pduReiAcc**, **!S:RConf**), (**!S:DReq**, **pduDisInd**), (**?S:CReq**, **pduConReq**), (**pduDisInd**, **!S:DInd**), (**?S:CReq**, **pduConReq**), (**pduConAcc**, **!S:CConf**), and (**pduDisInd**, **!S:DInd**)).

In order to allow the synthesized protocol to react to all possible messages in all states, random non-expected inputs are added to the list like ((**CReq**, **pduConAcc**), (**CReq**, **NULL**), (**pduConAcc**, **CConf**), ...). Such a list is then encoded and used as test set during the individual evaluation. For the evolution, de Araújo et al. [859] use the grammar-guided Genetic Programming Kernel (GPK, see Section 4.5.4 on page 149) of Hörner [991] which allows for variable-length genomes. The objective function (subject to maximization) steering how the protocols are grown compares the behavior of a solution candidate with the expectations and counts the number of correct outputs.

⁹ http://en.wikipedia.org/wiki/Protocol_data_unit [accessed 2008-06-21]




Fig. 23.9: The best individual from the experiment of de Araújo et al. [859].

Figure 23.9 illustrates the state-based representation (SBR) encoding of finite state machines as used by de Araújo et al. [859] on the example of the best individual resulting from their experiments. The genotypes are divided into genes, each representing one state. The total number of states is variable. Each gene contains a tuple of two numbers for each possible input. The first number in this tuple represents the next state and the second one stands for the output of the corresponding transition.

The receiver side of the protocol was evolved in the same manner during the experiments of de Araújo et al. [859] and the results were able to pass an automated equivalence check. The good performance reported further indicates the high utility of this approach for protocol synthesis.

Evolving finite state machine-based Protocol Definitions

Another approach for evolving finite state machine-based protocol specifications with genetic algorithms has been contributed by Sharples and Wakeman [710]. In [1648], for instance, protocols for unreliable media were evolved. In his PhD thesis, Sharples [711] shows that even complex protocols can be evolved that may even outperform protocols designed by hand. In most cases though, the fitness of the evolved protocols is lower.

Evolving Fraglet-based Protocols

In Section 4.9.2 on page 183, we have outlined the Fraglet language. According to Tschudin [1092], this form of protocol representation is predestined for automated synthesis via evolutionary algorithms: Fraglets have almost no syntactical constraints and can represent complicated protocols in a compact manner. Tschudin [1092] focused on the offline evolution

of protocols using a genetic algorithm. Therefore, a complete communication system was simulated for a given number of time steps during the evaluation of each individual. The objective values then denote the correlation of the behavior observed during the simulation and the target behavior. While Tschudin's results substantiate that the evolutionary methods are suitable to optimize existing Fraglet protocols, they also indicated that the evolution of new distributed algorithms is hard because of the all-or-nothing problem known from Genetic Programming (see Section 4.10.3 on page 188) and a strong tendency to overfitting (see Section 1.4.6 on page 48).

Online Protocol Adaptation and Evolution with Fraglets

Autonomic networks are networks where manual management is not desired or hard to realize, such as systems of hundreds of gadgets in an e-home, sensor networks, or arbitrary mesh networks with wireless and wired links. Yamamoto and Tschudin [707] propose that software in such networks needs to be self-modifying so as to be able to react to unforeseen network situations. They distinguish two forms of reaction – adaptation and evolution. Adaptation is the short-term reconfiguration of existing modules whereas evolution stands for the modification of old and the discovery of new functionality and happens at a larger timescale. Software with these abilities probably cannot predict whether the effects of a modification are positive or negative in advance and therefore needs to be resilient in order to mitigate faulty code that could evolve. In [1649], Tschudin and Yamamoto show that such resilience can be achieved by introducing redundancy into Fraglet protocols.

Complementing Tschudin's work [1092] on offline optimization summarized in Section 23.1.3, Yamamoto and Tschudin [707] describe online protocol evolution as a continuously ongoing, decentralized and asynchronous process of constant competition and selection of the most feasible modules. Genetic Programming with mutation and homologous crossover is chosen mechanism for accomplishing these features. The fitness measure (subject to maximization) is the performance of the protocols as perceived by the applications running in the network. The score of an individual (protocol) is incremented if it behaves correctly and decremented whenever an error is detected. The resource consumption in terms of the memory consumed by the protocols is penalized proportionally.

Yamamoto and Tschudin [1650; 1651] create populations containing a mix of different confirmed delivery and reliable delivery protocols for messages. These populations are then confronted with either reliable or unreliable transmission challenge and were able to adapt to these conditions quickly. If the environment changes afterwards, i. e., a formerly reliable channel becomes unreliable, for instance, the degree of re-adaptation is unsatisfying. The loss of diversity due to the selection of only highly fit protocols during the adaptation phase could not be compensated by mutation in these first experiments of Yamamoto and Tschudin [1650].

Further information on approaches for evolutionary online optimization of communication protocols can be found in the report *Framework for Distributed On-line Evolution of Protocols and Services, 2nd Edition* from the EU-sponsored project BIONETS [1652].

Learning Agent Communication Patterns

Some very early experiments in evolving communication patterns for autonomous agents were conducted by Werner and Dyer [891] on the example of artificial creatures. In their simulation, blind male animals capable of hearing and moving had to be guided by immobile females that were able to see and to produce sounds. A genetic algorithm was used to build neural networks for both, the females and the males.

Based on their co-evolutionary strategy [885], Iba et al. [886] evolve communication protocols used by agents for cooperative predator-prey pursuit and robot navigation tasks [887].

Mackin and Tazaki [888; 889; 890] focus on evolving communication patterns for agent in a multi-agent system modeling a market place. They use a divide-and-conquer approach for evolving the code for sending and receiving the messages exchanged amongst the agents in independent ADFs (automatically defined functions, see Section 4.3.9 on page 137) with Genetic Programming. While starting with a scenario where each trading agent has only a single objective in [888], the evolutionary approach is extended to multi-objective Genetic Programming. By building separate communication trees for each objective, Mackin and Tazaki [889] allow the agents to weight three objectives in their experiments: to buy at the highest volume at lowest price and to collect the best performing services for buying agents and to sell the highest volume at the highest price while collecting the highest value return for selling agents.

23.1.4 Optimizing Network Security

Global optimization can also be used for improving the security of networks.

Intrusion Detection

According to Heady et al. [1653], an intrusion is “any set of actions that attempt to compromise the integrity, confidentiality or availability of a resource”. Intrusion detection systems¹⁰ (IDS) identify such unwanted manipulation of computer systems that, most likely, come from the outside through network connections like the internet. Generally, a network connection can be described by a broad set of features, concerning its content, time- or traffic-based details, and host-related data. Analysis of these features boils down to data mining. Global optimization methods can be utilized for improving the efficiency of classifiers that decide whether a dataset containing these features belongs to normal or abnormal (possible intrusive) traffic.

Lu and Traore [864] and Heady et al. [1653], for instance, evolve classification rules intrusion detection rules with tree-based Genetic Programming. A similar approach is followed by Folino et al. [1654], who apply the AdaBoost ensemble method of Freund and Schapire [1583] in conjunction with the distributed evolution of decision trees. Linear Genetic Programming¹¹ also seems to be very efficient in growing such intrusion detectors, as shown by Mukkamala et al. [1655], for example. Song, Heywood, and Zincir-Heywood [865], [866] use their page-based LGP variant [1031] for this purpose. In [867], they extend this approach for identifying attacks on 802.11 WLAN networks. Hansen et al. [1656] apply a homologous¹² crossover operation to the evolution of machine-code linear structures on for intrusion detection for the prevention of cyber-terrorism.

Security Protocol Generation

Perrig and Song [1452] focus their work on the automatic generation of security protocols such as mutual authentication protocols. In the first step of this process, they use an iterative deepening depth-first search (IDDFS) algorithm (as introduced in Chapter 16 on page 245) to list *all* protocols that can be composed using a predefined set of primitives. The limit for this search is not the depth in the search space, but a value of a specified cost function which sums up user-specified costs of the single primitives of a protocol. Since the space of all protocols is very large, simplifying assumptions like the absence of redundant message fields and commutativity of the fields are made. Perrig and Song [1452] conduct experiments with the goal to synthesize mutual authentication protocols. In these experiments, the IDDFS still

¹⁰ http://en.wikipedia.org/wiki/Intrusion_detection_system [accessed 2008-06-16]

¹¹ You can find a discussion on linear Genetic Programming in Section 4.6 on page 161.

¹² Homologous crossover in linear Genetic Programming is discussed in Section 4.6.1.

generates tens of thousands of individuals even with the restrictions defined above. Therefore, a pruning algorithm is added which removes obviously flawed protocols such as those which openly disclose private keys in a messages. In this pruning process, also vulnerabilities to impersonation attempts or replay attacks are checked. Only a few hundreds of candidate protocols survive this pruning routine in the experiments mentioned. These candidates may still be incorrect and are therefore screened with Song's Athena security protocol checker [1657]. As a result, two correct symmetric and one correct asymmetric protocol were synthesized, whereas the symmetric one was even simpler than the *Symmetric-Key Three-Pass Mutual Authentication Protocol* standardized 1993 by the ISO [1658]. In 2001, Song et al. [1451] add a code generator that transforms the protocol specifications found to Java to their system.

23.1.5 Optimizing Parameters and Configurations

Protocols, distributed algorithms, and topologies are optimized in order to create robust and efficient distributed applications. The efficiency of such application, however, also strongly depends on its parameter setting and configuration. Today's web servers, for instance, have a variety of tuning parameters. It is hard to find optimal settings, especially when applications consist of web frontends, application servers hosting the business logic, and data base servers in the backend. The configurations of these subsystems are not independent. Instead, there are epistatic and synergetic effects which are often highly complicated and hard to simulate offline.

Optimal Application Server Configuration

Xi et al. [1277] have developed a smart hill climbing algorithm for application server configuration which needs few online performance evaluations of the running system for fine-tuning its parameters.

Distributed Agent Behavior Evolution

In the work of Nakano and Suda [483; 482; 1659], network services are represented by mobile agents. An agent capable of behaviors like replication, migration, and termination could, for instance, stand for a HTTP server. A behavior i is triggered when a corresponding weighted sum surpasses a certain threshold θ_i [483]. This sum incorporates environmental information such as the request rate and the resource costs and internal state of the agent like its age and energy level. The weights of the factors in this sum are subject to optimization.

Agents are furthermore equipped with an energy value. They obtain energy for performing services for the user and spend energy for resource access and replication. If the energy of agent is depleted, the agent dies, i. e., its service is terminated. Agents may create new agents by replication or reproduction. Replication corresponds to mutation in evolutionary algorithms, leading to a child with slightly changed weight vectors. Reproduction is basically a crossover operation where a partner is selected within n hops in the network based on three measures: the time a user has to wait between requesting and receiving a service, the number of hops a service travels from a user to agent, and the energy efficiency.

With their experiments, Nakano and Suda [483] showed that their approach allows the agents, i. e., the services, to adapt changing conditions such as network failures, varying workload, and changing platform (energy) costs.

Protocol Configuration

In his master's thesis, Grace [1660] uses Genetic Programming to evolve configurations for the JavaGroups [1661] reliable group communication framework. The behavior of a JavaGroups instance depends on a protocol stack specification in form of a string which is passed

in when it is created. This specification determines which protocol layers will be present in the stack. Grace [1660] showed that the evolution of communication protocol configurations for predefined scenarios, such as reliable, ordered multicast, is feasible with his approach.

Evolving Protocol Adapters

In distributed systems with components manufactured by different organizations, inter-process communication¹³ (IPC) can become problematic. Therefore, protocol adapters, components that mediate between different protocols, are applied. In the worst case, the number of required adaptors rises quadratically with the number of processes in the system. Hence, creating them automatically seems to be a good choice. Van Belle et al. [1662] evolve such protocol adapters in form of classifier systems representing Petri nets with a genetic algorithm. They show the viability of this method on the example of an adapter for a locking protocol for process synchronization.

Like me myself, they also found that evolving components of distributed systems always requires special means [1662] of overfitting prevention in order to stop the system from “cheating”. A controversial discussion of the use of Learning Classifier System for the purpose of protocol adapter synthesis can be found in Van Belle’s PhD thesis [1663].

23.2 Genetic Programming of Distributed Algorithms

23.2.1 Introduction

Distributed systems are one of the most vital components of our economy. While many internet technologies, protocols, and applications grew into maturity and have widely been researched, new forms of networks and distributed computing have emerged. Amongst them, we can find wireless networks [1664, 1665, 1666], sensor networks [1667] (and wireless sensor networks [1668]), Smart Home networks [1669, 1670], ubiquitous computing [1671, 1672], and ideas like amorphous computing [1673, 1674]. These distributed systems introduce new requirements like self-adaptation to change in the environment (nodes may enter and leave the networks frequently) or change the priority of others (such as energy consumption).

It may be a bold statement to say that such new requirements ask for new programming paradigms and future will shows whether it holds or not. Nobody will, however, argue that developing applications for these new distributed systems is surely to become more complicated than in traditional networks. Hence, exploring the utility of new programming methodologies (and new representations for algorithms especially tailored to them) is a demand of the current situation.

The design of a distributed algorithm is basically the transformation of a specification of the behavior of a network on the global scale to a program that must be executed locally on each of the nodes of the network in order to achieve this behavior. Up to now, no general method for automating this process illustrated in Fig. 23.10.1 has been developed and it is unlikely that this will change in near future.

The transformation of global system behavior to local rules is no process specific only to distributed algorithm design. Matter of fact, a widely studied example for are swarming behaviors in nature [1675, 1410]. These behaviors have evolved for millions of years. By allowing many individuals of a species to travel together in a configuration which has a good volume/surface-ratio, they improve defense against predators, increase the chance of finding mating partners, enhance foraging success, improve hydro or aerodynamics and so on. Nature has evolved many efficient swarming behaviors, such as the shoaling of fish (depicted in Fig. 23.10.2), flocking of birds, herding of cows, and swarming of locusts.

¹³ http://en.wikipedia.org/wiki/Inter-process_communication [accessed 2008-06-23]




Fig. 23.10.1: Design of distributed algorithms.




Fig. 23.10.2: Evolutionary behavior design.

Fig. 23.10: Global $\rightarrow$ Local behavior transformations.

Evolutionary algorithms copy the evolutionary process itself for solving complex optimization problems [409, 139] and Genetic Programming is the family of EAs which can be used for deriving programs [65]. Here we will utilize it for breeding distributed algorithms – in other words – for transforming descriptions of global behaviors to local algorithms.

These global descriptions are therefore encoded in objective functions, which rate “how close” the behavior of an evolved program x comes to the wanted one. In order to approximate its quality, we execute x on nodes represented by virtual machines in a whole simulated network. As in reality, many of these VMs run asynchronously at approximately the same speed, which may differ from VM to VM and cannot be assumed to be constant either. For different problems, different topologies are simulated.

We apply multi-objective Genetic Programming since it allows us to optimize the algorithms for different aspects during the evolution. While the functional objective functions perform the actual comparison of the observed behavior of the simulated network (running the evolved algorithms) with the desired global behavior, non-functional objective functions foster the economical use of resources, minimizing communication and program size, for instance.

23.2.2 Evolving Proactive Aggregation Protocols

In this section we discuss what proactive aggregation protocols are and how we can evolve them using a modified symbolic regression approach with Genetic Programming.

Aggregation Protocols

Definition 23.1 (Aggregate). In computer science, an aggregate function¹⁴ $\alpha : \mathbb{R}^m \mapsto \mathbb{R}$ computes a single result $\alpha(\mathbf{x})$ from a set of input data $\mathbf{x}$. This result represents some features of the input like its arithmetic mean.

Other examples for aggregate functions are the variance and the number of points in the input data. In general, an aggregate¹⁵ is a fusion of a (large) set of low-level data to one piece of high-level information. Aggregation operations in databases and knowledge bases [1676, 1677, 1678, 1679, 1680, 1681, 1682], be they local or distributed, for instance, have been an active research area in the past decades. Here, large datasets from different tables are combined to an aggregate by structured queries which need to be optimized for maximal performance.

With the arising interest in peer-to-peer applications (see Section 29.2.2) and sensor networks (discussed in Section 29.2.2), a whole new type of aggregation came into existence in the form of aggregation protocols. These protocols are a key functional building block by providing the processes in such distributed systems with access to global information including network size, average load, mean uptime, location and description of hotspots, and so on [1683, 1684]. Robust and adaptive applications often require this local knowledge of such properties of the whole. If, for example, the average concentration of some toxin (which is aggregated from the measurements of multiple sensors in a chemical laboratory) exceeds a certain limit, an alarm should be triggered.

In aggregation protocols, the data vector $\mathbf{x}$ is no longer locally available but its elements are spread all over the network. When computing the aggregate under these circumstances, we cannot just evaluate α . Instead, some form of data exchange must be performed by the nodes. This exchange can happen in two ways: either *reactive* or *proactive*. In a reactive aggregation protocol, one of the nodes in the network issues a query to all other nodes. Only this node receives the answer in form of the result (the aggregate) or the data needed to compute the result as illustrated in Fig. 23.11.1. A proactive aggregation protocol, as sketched in Fig. 23.11.2, on the other hand allows all nodes in the network to receive knowledge of the aggregate. This is achieved by repetitive data exchange amongst the nodes and iterative local refinement of the estimates of the wanted value. Notice that the trivial solution would be that all nodes send their information to all other nodes. Generally, this is avoided since it is not a viable approach and instead, the data is disseminated step by step as part of the estimates.

Gossip-Based Aggregation

Jelasity, Montresor, and Babaoğlu [1684] propose a simple yet efficient type of proactive aggregation protocols [1685]. In their model, a network consists of many nodes in a dynamic topology where every node can potentially communicate with every other node. Errors in communication may occur, Byzantine faults not. The basic assumption of the protocol is that each node in the network holds one numerical value x . This value represents some information about the node or its environment, like, for example, the current work load. The task of the protocol is to provide all nodes in the network with an up-to-date estimate of the aggregate function $\alpha(\mathbf{x})$ of the vector of all values $\mathbf{x} = (x_p, x_q, \dots)^T$.

The nodes hold local states s (possibly containing x) which they can exchange via communication. Therefore, each node knows picks its communication partners with the `getNeighbor()` method.

The skeleton of the gossip-based aggregation protocol is specified in Algorithm 23.1 and consists of an active and a passive part. Once in each $\delta > 0$ time units, at a randomly picked time, the active thread of a node p selects a neighbor q . Both partners exchange their

¹⁴ http://en.wikipedia.org/wiki/Aggregate_function [accessed 2007-07-03]

¹⁵ http://en.wikipedia.org/wiki/Aggregate_data [accessed 2007-07-03]




Fig. 23.11.1: reactive aggregation




Fig. 23.11.2: proactive aggregation

Fig. 23.11: The two basic forms of aggregation protocols.

information and update their states with the update method: p calls $\text{update}(s_p, s_q)$ in its active thread and q calls $\text{update}(s_q, s_p)$ in the passive thread. update is defined according to the aggregate that we want to be computed

Algorithm 23.1: gossipBasedAggregation()

Data: p : the node running the algorithm

Data: s_p : the local state of the node p

Data: s_q, s_r : states received as messages from the nodes q and r

Data: q, p, r : neighboring nodes in the network

```

1 begin
2   Subalgorithm activeThread
3   begin
4     while true do
5       do exactly once in every  $\delta$  units at a randomly picked time:
6        $q \leftarrow \text{getNeighbor}()$ 
7        $\text{sendTo}(q, s_p)$ 
8        $s_q \leftarrow \text{receiveFrom}(q)$ 
9        $s_p \leftarrow \text{update}(s_p, s_q)$ 
10    end
11   Subalgorithm passiveThread
12   begin
13     while true do
14        $s_r \leftarrow \text{receiveAny}()$ 
15        $\text{sendTo}(\text{getSender}(s_r), s_p)$ 
16        $s_p \leftarrow \text{update}(s_p, s_r)$ 
17    end
18 end

```

Example – Distributed Average

Assume that we have built a sensor network measuring the temperature as illustrated in Figure 23.12. Each of our sensor nodes is equipped with a little display visible to the public.




Fig. 23.12: An example sensor network measuring the temperature.

The temperatures measured locally will fluctuate because of wind or light changes. Thus, the displays should not only show the temperature measured by the sensor node they are directly attached to, but also the average of all temperatures measured by all nodes. Then, the network needs to execute a distributed aggregation protocol in order to estimate that average.

If we therefore choose a gossip-based average protocol, each node will hold a state variable which contains its local estimation of the mean. The update function, henceforth receiving the local approximation and the estimate of another node, returns the mean of its inputs.

$$\text{update}_{\text{avg}}(s_p, s_q) = \frac{s_p + s_q}{2} \quad (23.5)$$

If two nodes p and q communicate with each other, the new value of s_p and s_q will be $s_p(t+1) = s_q(t+1) = 0.5 * (s_p(t) + s_q(t))$. The sum – and thus also the mean – of both states remains constant. Their variance, however, becomes 0 and so the overall variance in the network gradually decreases.




Fig. 23.13.1: initial state




Fig. 23.13.2: after step 1




Fig. 23.13.3: after step 2

Fig. 23.13: An gossip-based aggregation of the average example.

In order to visualize how that type of protocol works, let us assume that we have a network of four nodes with the initial values $\mathbf{x} = (5, 6, 7, 8)^T$ as illustrated in Fig. 23.13.1. The arithmetic mean of its elements is

$$\frac{5 + 6 + 7 + 8}{4} = \frac{13}{2} = 6.5 \quad (23.6)$$

The initial variance is

$$\frac{(5 - 6.5)^2 + (6 - 6.5)^2 + (7 - 6.5)^2 + (8 - 6.5)^2}{4} = \frac{5}{4} \quad (23.7)$$

In the first step of the protocol, the nodes with the initial values 5 and 7 as well as the other two exchange data with each other and update their values to 6 and 7 respectively (see Fig. 23.13.2). Now the average of all estimates is still

$$\frac{6 + 6 + 7 + 7}{4} = 6.5 \quad (23.8)$$

but the variance has been reduced to

$$\frac{(6 - 6.5)^2 + (6 - 6.5)^2 + (7 - 6.5)^2 + (7 - 6.5)^2}{4} = 1 \quad (23.9)$$

After the second protocol step, outlined in Fig. 23.13.3, all nodes estimate the mean with the correct value 6.5 (and thus, the variance is 0). The distributed average protocol is only one example of gossip-based aggregation. Others are:

- **Minimum and Maximum.** The minimum and maximum of a value in the network can be computed by setting $\text{update}_{\min}(s_p, s_q) = \min\{s_p, s_q\}$ and $\text{update}_{\max}(s_p, s_q) = \max\{s_p, s_q\}$ respectively.
- **Count.** The number of nodes in a network N can be computed using the average protocol: the initiator sets its state to 1 and all other nodes begin with 0. Then the average is computed is then $\frac{1+0+0+\dots}{\text{len}(N)} = \frac{1}{\text{len}(N)}$ where $\text{len}(N)$ is the number of nodes in N . The nodes now just need to invert the computed value locally and obtain $\frac{1}{\frac{1}{\text{len}(N)}} = \text{len}(N)$.
- **Sum.** The sum of all values in the network can be computed by estimating both, the mean value $\bar{x}$ and the number of nodes $\text{len}(N)$ in the network N simultaneously and multiplying both with each other: $\text{len}(N) \bar{x} = \sum x$.
- **Variance.** As declared in Equation 27.61 on page 424, the variance of a data set is the difference of the mean of the squares of the values and the square of their means. Therefore, if we compute $\bar{x^2}$ and $\bar{x}$ by using the average protocol, we can subtract them $\text{var}(x) \approx \bar{x^2} - \bar{x}^2$ and, hence, obtain an estimation of the variance.

Further considerations are required if $\mathbf{x}$ is not constant but changes by and by. Both, peer-to-peer networks as well as sensor networks, have properties (discussed in Section 29.2.2) which are very challenging for distributed applications and lead to an inherent volatility of $\mathbf{x}$. According to Jelasity et al. [1684], a default approach to handle unstable data is to periodically restart the aggregation protocols. In our research, we were able to provide alternative aggregation protocols capable of dealing with dynamically changing data. This approach is discussed in Section 23.2.2 on page 366.

The Solution Approach: Genetic Programming

In order to derive certain aggregate functions automatically, we could modify the Genetic Programming approach for symbolic regression introduced in Section 22.1 on page 341 [863, 1686]. Let $\alpha : \mathbb{R}^m \mapsto \mathbb{R}$ be the exact aggregate function. It works on a vector of the dimension m containing the data elements. $m \in \mathbb{N}$ is not a predetermined constant but depends on the network size, i. e., $m = \text{len}(N)$ and α will return exact results for $m = 1, 2, 3, \dots$

In Section 27.6.1 on page 453, we will show that the dimension m of the domain $\mathbb{R}^m$ of α plays no role when approximating it with a maximum likelihood estimator. The theorems used there are again applied in symbolic regression (see Equation 22.6 on page 343), so the value of m does not affect the correctness of the approach. Deriving aggregation functions for distributed systems, however, exceeds the capabilities of normal symbolic regression. Each of the $m = \text{len}(N)$ nodes in the network N holds exactly one element of the data vector. α cannot be computed directly anymore since it requires access to all data elements at once. Instead, each node has to execute local rules that define how data is exchanged and how an approximation of the aggregate value is calculated. How to find these rules automatically is the subject of our research here. There are three use cases for such an automated aggregation protocol generation:

- We may already know a valid aggregation protocol but want to find an equivalent protocol which has advantages like faster convergence or robustness in terms of input volatility. This case is analogous to finding arithmetic identities in symbolic regression.
- We do not know the aggregate function α nor the protocol but have a set of sample data vectors $\mathbf{x}_i$ (maybe differing in dimensionality) and corresponding aggregates y_i . Using Genetic Programming, we attempt to find an aggregation protocol that fits to this sample information.
- The most probable use case is that we know how to compute the aggregate locally with a given function α but want to find a distributed protocol that does the same. We, for example, are well aware of how to compute the arithmetic mean of a data set $(x_1, x_2, \dots, x_m)$ – we just divide the sum of the single data items by their number m . If these items, however, are distributed and not locally available, we cannot simply sum them up. The correct solution described in Section 23.2.2 on page 368 is that each node starts by approximating the mean with its locally known value. Now always two nodes inform each other about their estimates and set their new approximation to be that mean of the old and the received one. This way, the aggregate is approached by iteratively refining the estimations.
The transformation of the local aggregate calculation rule α to the distributed one is not obvious. Instead of doing it by hand, we can just use the local rule to create sample data sets and then apply the approach of the second use case.

Network Model and Simulation

For gossip-based aggregation protocols, Jelasity et al. [1684] assume a topology where all nodes can potentially communicate with each other. In this fully connected overlay network, communication can be regarded as fault-free. Taking a look at the basic algorithm scheme of such protocols introduced as Algorithm 23.1 on page 368, we see that the data exchange happens once every δ time units at a randomly picked point in time. Even though being asynchronous in reality, it will definitely happen in this time span. That is, we may simplify the model to a synchronous network model where all communication happens simultaneously.

Another aspect of communication is how the nodes select their partners for the data exchange. It is a simple fact that the protocol can only converge to the correct value if each node has, maybe over multiple hops and calculations, been able to receive information from all other nodes. Imagine a network N consisting of $m = \text{len}(N) = 4$ nodes p , q , r , and t . If the communication partners are always (p, q) and (r, t) , the data dissemination is insufficient since p will never be able to incorporate the knowledge of the states of r and t . On the other hand, one data exchange between q and r will allow the protocol to work since p would later on indirectly receive the required information from q .

Besides this basic fact, Jelasity, Montresor, and Babaoğlu [1684] have shown that different forms of pair selection influence the convergence speed of the protocol. Correct protocols will always converge if complete data dissemination is guaranteed. Knowing that, we should choose a partner selection method that leads to fast convergence because we then can safe

protocol steps in the evaluation process. The pair building should be deterministic, because randomized selection schemes lead to slow convergence [1684], and, more importantly, will produce different outcomes in each test and make comparing the different evolved protocols complicated (as discussed in Section 1.3.4 on page 36). Therefore, choosing a deterministic selection scheme seems to be the best approach. Perfect matching according to Jelasity et al. [1684] means that each node is present in exactly one pair per protocol cycle, i. e., always takes part in the data exchange. If different pairs are selected in each cycle, the convergence speed will increase. It can further be increased by selecting (different) pairs in a way that disseminates the data fastest.




Fig. 23.14.1: pair-based




Fig. 23.14.2: general

Fig. 23.14: Optimal data dissemination strategies.

From these ideas, we can derive a deterministic pair selection mechanism with best-case convergence. Therefore, we first need to set the number of nodes in the simulated network $\text{len}(\mathbf{N}) = m = 2^d$ as a power of two. In each protocol step t with $t = 1, 2, \dots$, we compute a value $\Delta = 2^{t \bmod d}$. Then we build pairs in the form $(i, i + \Delta)$, where i and $i + \Delta$ are the indices identifying the nodes. This setup equals a butterfly graph and is optimal, as you can see in Fig. 23.14.1. The data from node 0 (marked with a thick border) spreads in the first step to node 1. In the second step, it reaches node 2 directly and node 3 indirectly through node 1. Remember, if the average protocol would use this pair selection scheme, node 3 would compute its new estimate at step 2 since

$$s_3(t=2) = \frac{s_3(t=1) + s_1(t=1)}{2} + \frac{\frac{s_3(t=0) + s_2(t=0)}{2} + \frac{s_0(t=0) + s_1(t=0)}{2}}{2} \quad (23.10)$$

In the third protocol step, the remaining four nodes receive knowledge of the information from node 0 and the data is disseminated over the complete network. Now the cycle would start over again and node 0 would communicate with node 1.

This pair selection method is bounded to networks of the size $m = 2^d$. We can generalize this approach by breaking up the strict pair-communication limitation. Therefore, we set $d = \lfloor \log_2 m \rfloor$ while still leaving $\Delta = 2^{t \bmod d}$ and define that a node i sends its data to the node $(i + \Delta) \bmod m$ for all i as illustrated in Fig. 23.14.2. This general communication rule abandons the strict pair-based data exchange but leaves any other feature of the aggregation protocols, like the working of the update method, untouched. We should again visualize that this rule is only defined so we can construct simulations where the protocols need as few

as possible steps to converge to the correct value in order to spare us computation time. Another important aspect also becomes obvious here: The time that an aggregation protocol needs to converge will always depend on the number of nodes in the (simulated) network.

Node Model and Simulation

As important as modeling the network is the model of the nodes it consists of. In Figure 23.15, we illustrate an abstraction especially suitable for fast simulation of aggregation protocols. A node p executing a gossip-based aggregation protocol receives input in form of the locally




Fig. 23.15: The model of a node capable to execute a proactive aggregation protocol.

known value (for example a sensor reading) and also in form of messages containing data from other nodes in the network. The output of p consists of the local approximation of the aggregate value and the information sent to its partners in the network. The computation is done by a processor which updates the local state by executing the update function. The local state s_p of p can most generally be represented as vector $\mathbf{s}_p \in \mathbb{R}^n$ of the dimension n , where n is the number of memory cells available on a node.

Like Jelasity et al. [1684], we until now have considered the states to be scalars. Generalizing them to vectors allows us to specify or evolve more complicated protocols. The state vector contains the approximation of the aggregate value at the position $i : 1 \leq i \leq n$. If the state only consists of a single number, the messages between two nodes will always include this single number and hence, the complete state.

A state vector not only serves as a container for the aggregate, but also as memory capable of accumulating information. It is probably unnecessary or unwanted to exchange the complete state during the communication. Therefore, we specify an index list e containing the indices of the elements to be sent and a list r with the indices of the elements that shall receive the values of the incoming messages. For a proper communication between the nodes, the length of e and r must be equal and each index must occur at most once in e and also at most once in r . Whenever a node p receives a message from node q , the following assignment will be done, with $\mathbf{s}_p[i]$ being the i^{th} component of the vector:

$$\mathbf{s}_p[r_j] \leftarrow \mathbf{s}_q[e_j] \quad \forall j = 0 \dots \text{len}(r) - 1 \quad (23.11)$$

In the original form of gossip-based aggregation protocols, the state is initialized with a static input value which is stepwise refined to approximate the aggregate value [1684]. In our

model, this restriction is no longer required. We specify an index I pointing at the element of the state vector that will receive the input. This allows us to grow protocols for static and for volatile input data. In the latter case, the inputs are refreshed in each protocol step. A node p would then perform

$$\mathbf{s}_p[I](t) \leftarrow \text{getInput}(p, t) \quad (23.12)$$

The function $\text{getInput}(p, t)$ returns the input value of node p at time step t . With this definition, the state vectors $\mathbf{s}_p$ become time-dependent, written as $\mathbf{s}_p(t)$. Finally, update is now designed as a map $\mathbb{R}^n \mapsto \mathbb{R}^n$ to return the new state vector.

$$\mathbf{s}_p(t+1) = \text{update}(\mathbf{s}_p(t)) \quad (23.13)$$

In the network simulation, we can put the state vectors of all nodes together to a single $n \times m$ matrix $S(t)$. The column k of this matrix contains the state vector $\mathbf{s}_k$ of the node k .

$$S(t) = (\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_m) \quad (23.14)$$

$$S[j, k] = \mathbf{s}_k[j] \quad (23.15)$$

This notation is used in Algorithm 23.2 and Algorithm 23.3. In Algorithm 23.2 we specify how the model definitions just discussed can be used to build a network simulation for gossip-based, proactive aggregation protocols. Here we also apply the general optimal communication scheme explained in Section 23.2.2. In the practical realization, we can spare creating a new matrix $S(t)$ in each time step t by initial using two matrices S_1, S_2 which we simple swap in each turn.

Evaluation and Objective Values

The models described before are the basis of the evaluation of the aggregation protocols that we breed. In general, there are two functional features that we want to develop in the artificial evolution:

1. We want to grow aggregation protocols where the deviation between the local estimates and the global aggregate is as small as possible, ideally 0.
2. This deviation can surely not be 0 after the first iteration at $t = 1$, because the nodes do not know all data at that time. However, the way how received data is incorporated into the local state of a node can very well influence the speed of convergence to the wanted value. Therefore, we want to find protocols that converge as quickly as possible.

In all use cases discussed in Section 23.2.2, we either already know the correct aggregation values y_i or the local aggregate function $\alpha : \mathbb{R}^m \mapsto \mathbb{R}$ that calculates them from data vectors of the length m . The objective is to find a distributed protocol that computes the same aggregates in a network where the data vector is distributed over m nodes. In our model, the estimates of the aggregate value can be found at the positions $S[\cdot, \cdot] \star \equiv \mathbf{s}_k[\cdot] \forall k \in [1 \dots m]$ in the state matrix or the state vectors respectively.

The deviation $\varepsilon(k, t)$ of the local approximation of a node k from the correct aggregate value $y(t)$ at a point in time t denotes its estimation error.

$$y(t) = \alpha\left(\left(\text{getInput}(1, t), \dots, \text{getInput}(m, t)\right)^T\right) \quad (23.16)$$

$$\varepsilon(k, t) = y(t) - S[\cdot, k](t) = y(t) - \mathbf{s}_k[\cdot] \quad (23.17)$$

We have already argued that the mean square error is an appropriate quality function for symbolic regression (see Equation 22.6). Analogously, the mean of the squares of the errors ε over all simulated time steps and all simulated nodes is a good criterion for the utility of an

Algorithm 23.2: simulateNetwork(m, T)

Input: m : the number of nodes in the simulation
Input: T : the maximum number of simulation steps
Input: [implicit] $update$: the update function
Input: [implicit] I : the index for the input values
Input: [implicit] O : the index for the output values
Input: [implicit] e : the index list for the send values
Input: [implicit] r : the index list for the receive values
Data: d : communication step base according to Fig. 23.14.2 on page 372
Data: k : a node index
Data: $S(t)$: the simulation state matrix at time step t
Data: Δ : the communication partner offset
Data: p : the communication partner node

```

1 begin
2    $d \leftarrow \lceil \log_2 m \rceil$ 
3    $S(0) \leftarrow \text{createMatrix}(n \times m)$ 
4   // initialize with local values
5    $S(0)[\cdot, k] \leftarrow \text{getInput}k0$ 
6    $t \leftarrow 1$ 
7   while  $t \leq T$  do
8      $S(t) \leftarrow \text{copyMatrix}(S(t-1))$ 
9      $\Delta \leftarrow 2^{t \bmod d}$ 
10    // perform communication according to Fig. 23.14.2 on page 372
11     $k \leftarrow 1$ 
12    while  $k \leq m$  do
13       $p \leftarrow (k + \Delta) \bmod m$ 
14      foreach  $j \in [1..\text{len}(r)]$  do  $S(t)[r_j, p] \leftarrow S(t-1)[e_j, k]$ 
15       $k \leftarrow k + 1$ 
16    // set (possible) new input values and perform update
17     $k \leftarrow 1$ 
18    while  $k \leq m$  do
19       $S(t)[I, k] \leftarrow \text{getInput}(k, t)$ 
20       $S(t)[i_j, k] \leftarrow \text{update}(S(t)[\cdot, k])$ 
21      //  $\equiv s_k(t) \leftarrow \text{update}(s_k(t))$ 
22       $k \leftarrow k + 1$ 
23     $t \leftarrow t + 1$ 
24 end
  
```

aggregation protocol. It is even related to both functional aspects subject to optimization: The larger it is, the greater is the deviation of the estimates from the correct value. If the convergence speed of the protocol is low, these deviations will become smaller more slowly by time. Hence, the mean square error will also be higher. For any evolved update function u we define¹⁶:

$$f_1(u, e, r) = \frac{1}{T * m} \sum_{t=1}^T \sum_{k=1}^m \varepsilon(k, t)^2 \Big|_{u, e, r} \quad (23.18)$$

This rather mathematical definition is realized indirectly in Algorithm 23.3, which returns the value of f_1 for an evolved update method u . It also applies the fast, convergence-friendly communication scheme discussed in Section 23.2.2. Its realization in the Distributed Genetic Programming Framework [862] software allows us to evaluate even complex distributed protocols in very short time: A protocol can be tested on 16 nodes for 300 protocol steps less than 5 milliseconds on a normal, 3GHz off-the-shelf PC.

¹⁶ where $\cdot|_{u, e, r}$ means “passing u, e, r as input to Algorithm 23.3”

Algorithm 23.3: $f_1(u, e, r) \leftarrow \text{evaluateAggregationProtocol}(u, m, T)$

Input: u : the evolved protocol update function to be evaluated
Input: m : the number of nodes in the simulation
Input: T : the maximum number of simulation steps
Input: [implicit] update: the update function
Input: [implicit] I : the index for the input values
Input: [implicit] O : the index for the output values
Input: [implicit] e : the index list for the send values
Input: [implicit] r : the index list for the receive values
Data: d : communication step base according to Fig. 23.14.2 on page 372
Data: k : a node index
Data: $S(t)$: the simulation state matrix at time step t
Data: Δ : the communication partner offset
Data: p : the communication partner node
Data: res : the variable accumulating the square errors
Output: $f_1(u, e, r)$: the sum of all square errors (deviations from the correct aggregate) over all time steps

```

1 begin
2    $d \leftarrow \lceil \log_2 m \rceil$ 
3    $S(0) \leftarrow \text{createMatrix}(n \times m)$ 
4   // initialize with local values
5    $S(0)[\cdot, k] \leftarrow \text{getInput}k0$ 
6    $t \leftarrow 1$ 
7   while  $t \leq T$  do
8      $S(t) \leftarrow \text{copyMatrix}(S(t-1))$ 
9     ...
10    // perform communication according to Fig. 23.14.2 on page 372
11    ...
12    // set (possible) new input values and perform update
13     $k \leftarrow 1$ 
14    while  $k \leq m$  do
15       $S(t)[I, k] \leftarrow \text{getInput}(k, t)$ 
16      //  $u$  is the evolved update-function and thus, used here
17       $S(t)[\cdot, k] \leftarrow u(S(t)[\cdot, k])$ 
18       $res \leftarrow res + (y(t) - S(t)[O, k])^2$ 
19      //  $\equiv res \leftarrow res + (\alpha(i(t)) - S(t)[o, k])^2$ 
20       $k \leftarrow k + 1$ 
21     $t \leftarrow t + 1$ 
22  return  $res$ 
23 end
  
```

Input Data

In Algorithm 23.3 we use sample α values in order to determine the errors ε . In two of our initial use cases, we need to create these values before the evaluation process, either with an existing protocol or with a known aggregate function α . Here we will focus on the latter case.

If transforming a local aggregate function α to a distributed aggregation protocol, we need to create sample data vectors for the $\text{getInput}(k, t)$ -method. Here we can differentiate between *static* and *dynamic* input data: for static input data, we just need to create the samples for $t = 0$ since $\text{getInput}(k, 0) = \text{getInput}(k, 1) = \dots \text{agGetInput}bkT \forall k \in [1..n]$. If we have dynamic inputs on the other hand, we need to ensure that at least some elements of the input vectors $\mathbf{x}(t) = \text{getInput} \cdot t$ will differ, i. e., $\exists t_1, t_2 : \mathbf{x}(t_1) \neq \mathbf{x}(t_2)$. If this difference is too large, an aggregation protocol cannot converge. It should be noted that it would

be wrong to assume that we can measure this difference in terms of the sample data $\mathbf{x}$ – restrictions like $0.9 < \left| \frac{\mathbf{x}_{[i]}(t)}{\mathbf{x}_{[i]}(t+1)} \right| < 1.1$ are useless, because their impact on the value of α is unknown. Instead, we must limit the variations in terms of the aggregation results, like

$$0.9 < \left| \frac{\alpha(\mathbf{x}(t))}{\alpha(\mathbf{x}(t))} \right| < 1.1 \quad (23.19)$$

In both, the static and the dynamic case, we need to create multiple input datasets, distinguished by adding a dataset index to $\mathbf{x}(t)$: $\mathbf{x}(1, t), \mathbf{x}(2, t), \dots, \mathbf{x}(l, t)$. Only if $\alpha(\mathbf{x}(1, t)) \neq \alpha(\mathbf{x}(2, t)) \neq \dots \neq \alpha(\mathbf{x}(l, t))$ we can assure that the result are not just overfitted protocols that simple always return one single learned value: the exact α of the sample data. The single $\mathbf{x}(i, t)$ should differ in magnitude, sign, and distribution since this will lead to large differences in the α -values:

$$\left(\left| \frac{\alpha(\mathbf{x}(i, t))}{\alpha(\mathbf{x}(j, t))} \right| \ll 1 \right) \vee \left(\left| \frac{\alpha(\mathbf{x}(i, t))}{\alpha(\mathbf{x}(j, t))} \right| \gg 1 \right) \quad \forall i \neq j \quad (23.20)$$

We use z such data sets to perform z runs of Algorithm 23.3 and compute the true value of f_1 as arithmetic mean of the single results.

$$f_1(u, e, r) = \frac{1}{z} \sum_{i=1}^z \text{underCondition} f_f(u, e, r) \mathbf{x}_i \quad (23.21)$$

Of course, for each protocol that we evaluate we will use the same sample data sets because otherwise the results would not be comparable. It should be noted that overfitting can furthermore be prevented by changing the sample vectors in each generation. In this experiment series, generating the test data was too time consuming so we did not apply this measure.

Volatile Input Data

The specification of $\text{getInput}(k, t)$ which returns the input value of node k at time $t \in [0..T]$ allows us to evolve aggregation protocols for static and such for volatile input. Traditional aggregation protocols are only able to deal with constant inputs [1684]. These protocols have good convergence properties, as illustrated in Fig. 23.16.1. They always converge to the correct results but will simple ignore changes in the input data (see Fig. 23.16.2).

They would need to be restarted in a real application from time to time in order to provide up-to-date approximations of the aggregate. This approach is good if the input values in the real application that we evolve the protocols for change slowly. If these inputs are volatile, the estimations of these protocols become more and more imprecise. The fact that an aggregation protocol needs a certain number of cycles to converge is an issue especially in larger or mobile sensor networks. One way to solve this problem is to increase the data rate of the network accordingly and to restart the protocols more often. If this is not feasible, because of, for example, energy restrictions in a low-power sensor network application prohibit increasing the network traffic, dynamic aggregation protocols may help.

They represent a sliding average of the approximated parameter and are able to cope with changing input data. In each protocol step, they will incorporate their old state, the received information, and the current input data into the calculations. A dynamic distributed average protocol like the one illustrated in Figure 23.17 is a weighted sum of the old estimate, the received estimate, and the current value. The weights in the sum can be determined by the Genetic Programming process according to the speed with which the inputs change. In order to determine this speed for the simulations, a few real sample measurements would suffice to produce customized protocols for each application situation. However, the incorporation of the current input value is also the drawback of such an approach, since it cannot fully converge to the correct result anymore.




Fig. 23.16.1: with constant inputs




Fig. 23.16.2: with volatile inputs

Fig. 23.16: The behavior of the distributed average protocol in different scenarios.




Fig. 23.17: A dynamic aggregation protocol for the distributed average.

Phenotypic Representations of Aggregation Protocols

We have to find a proper representation for gossip-based aggregation protocols. Such a protocol consists of two parts: the evolved update function and a specification of the properties of the state vector – the variables I , O , r , and e .

Representation for the update Function

The function $update$ as defined in the context of our basic model for aggregation protocols receives the state vectors $\mathbf{s}_k(t) \in \mathbb{R}^n$ of time step t as input. It returns the new state vectors $\mathbf{s}_k(t+1) \in \mathbb{R}^n$ of time step $t+1$. This function is indeed an algorithm by itself which can be represented as a list of tuples $l = (\dots, (u_j, v_j), \dots)$ of mathematical expressions u_j and vector element indices v_j . This list l is processed sequentially for $j = 1, 2, \dots, \text{len}(l)$. In each step j , the result of the expression u_j is computed and assigned to the v_j^{th} element of the old state vector $\mathbf{s}(t-1)$. In the simplest case, l will have the length $\text{len}(l) = 1$. One example for this is the well-known distributed average protocol illustrated in Figure 23.18: In the single formula, the first element of $\mathbf{s}_{[1]}(t)$, $[1]$, is assigned to $0.5 * ([1] + [2])$ which is the average of its old value and the received information. Here, the value of the first element is sent to the partner and the received message is stored in the second element, i. e., $r = [2]$, $e = [1]$. The terminal set of the expressions now does not contain the simple variable x anymore but all elements of the state vectors. Finally, after all formulas in the list have been computed and their return values are assigned to the corresponding memory cells, the modified old state vector $\mathbf{s}_k(t)$ becomes the new one $\mathbf{s}_k(t+1)$. Fig. 23.18.2 shows a more complicated protocol where $update$ consists of $\text{len}(l) = 4$ formulas $((u_1, 1), (u_2, 2), (u_3, 3), (u_4, 2))$. We will not elaborate deeper on these examples but just note that both are valid results of Genetic Programming – a more elaborate discussion of them can be found in Section 23.2.2 on page 382, Section 23.2.2 on page 383, and [1686, 863].




Fig. 23.18.1: distributed average




Fig. 23.18.2: square root of the distributed average

Fig. 23.18: Some examples for the formula series part of aggregation protocols.

The important point is that we are able to provide a notation for the first part of the aggregation protocol specification that is compatible to normal symbolic regression and which thus can be evolved using standard operators.

Besides this sequence of formulas computed repetitively in a cycle, we also need an additional sequence that is executed only once, in the initialization phase. It is needed for some other protocols than the distributed minimum, maximum, and average, which cannot assume the approximation of the estimate to be the current input value. Here, another sequence of instructions is needed which transforms the input value into an estimate which then can be exchanged with other nodes and used as basis for subsequence calculations. This additional sequence is evolved and treated exactly in the same way as the set of formulas used inside the protocol cycle.

Experiments have shown that it is useful though to initialize all state elements in the first time step with the input values. Therefore, both Algorithm 23.2 and Algorithm 23.3, initially perform $S(0)_{[\cdot, k]} \leftarrow \text{getInput}(k, 0)$ instead of $S(0)_{[i, k]} \leftarrow \text{getInput}(k, 0)$. In all other time steps, only $S(t)_{[i, k]}$ is updated.

Straightforwardly, we can specify a non-functional objective function f_2 that returns the number of expressions in both sets and, hence, puts pressure into the direction of small protocols with less computational costs.

Representation for I , O , e , and r

Like the update function, the parameters of the data exchange, r and e , become subject to evolution. I and O are only single indices; we can assume them to be fixed as $I = 1$ and $O = 2$. Allowing them to be changed will only result in populations of many incompatible protocols. Although we could do the same with e and r , there is a very good reason to make them variable. If e and r are built during the evolutionary process, different protocols with different message lengths ($\text{len}(e_1) \neq \text{len}(e_2)$) can emerge. Hence, we can introduce a non-functional objective function f_3 that puts pressure into the direction of minimal message lengths. The results of Genetic Programming will thus be optimal not only in accuracy of the results but only in terms of communication costs.

For the lists e and r there are three possible representations. We can use either a bit string of the fixed length $2n$ which contains two bits for each element of s : the first bit determines if the value of the element should be sent, the second bit denotes if an incoming element should be stored there. String genomes of a fixed length are explained in detail in Section 3.4 on page 119. By doing so, we implicitly define some restrictions on the message structure since we need to define an order on the elements inside. If $n = 4$, a bit string 01011010 will be translated into $e = (3, 4)$ and $r = (1, 2)$. It is not possible to obtain something like $e = (3, 4)$ and $r = (2, 1)$.

The second encoding scheme is to use two variable-length integer strings which represent e and r directly. Such genomes are introduced in Section 3.5 on page 121. Now the latter

case becomes possible. If the lengths of the two strings differ, for example for reproduction reasons, the length of the shorter one is used solely.

The third approach would be to, again, evolve one single string z . This string is composed of pairs $z = ((e_1, r_1), (e_2, r_2), \dots, (r_l, r_l))$. The second and the third approach are somewhat equivalent,

In principle, all three methods are valid and correct since the impossibility of some message structures in the first method does not necessarily imply that certain protocol functionality cannot evolve. The standard reproduction operators for string genomes, be they of fixed or variable length, can be applied.

When we closely examine our abstract protocol representation, we will see that it will work with epidemic [1687] or SPIN-based [1688] communication too, although we developed it for a gossip-based communication model.

Reproduction Operators

As already pointed out when elaborating on the representation schemes for the two parts of the aggregation protocols, well-known reproduction operators can be reused here.

- The formulas in the protocol obey strictly a tree form, where the root always has two child nodes, the formula sequences for the protocol cycle and the initialization, which, in turn, may have arbitrarily many children: the formulas themselves. A formula is a tree node which has stored one number which identifies the vector element its results will be written to. It has exactly one child node, the mathematical expression which is a tree of other expressions. We elaborate on tree-shaped genomes in Section 4.3 on page 132.
- The communication behavior is described as either one fixed-length bit string or two variable-length integer strings.

New protocols are created by first building a new formula tree and then combining it with one (or two, according to the applied coding scheme) newly created string chromosomes. We define the mutation operation as follows: If an aggregation protocol is mutated, with 80% probability its formula tree is modified and with 20% probability its message pattern. When performing a recombination operation, a new protocol is constructed by recombining the formula tree as well as the message definition of both parents with the default means.

Results from Experiments

In Table 23.2, we list the configuration of the Genetic Programming algorithm applied to breed aggregation protocols.

Parameter	Short	Description
Problem Space	$\mathbb{X}$	The space of aggregation programs. (see Section 23.2.2)
Objective Functions	F	$F = \{f_1, f_2, f_3\}$, see Algorithm 23.3, Section 23.2.2, Section 23.2.2
Search Space	$\mathbb{G}$	(basically) identical with the problem space, i. e., $\mathbb{G} = \mathbb{X}$.
Search Operations	Op	mutation and crossover (see Section 23.2.2)
GPM	gpm	not needed
Optimization Algorithm	alg	elitist Genetic Programming
Comparison Operator	cm	$\text{cmp}_{F,agg}$ (see Equation 23.22)

Population Size	ps	$ps = 4096$
Maximum Archive Size	as	The size of the archive with the best known individuals was limited to $as = 64$. (see Definition 2.4)
Steady-State	ss	The algorithms were generational (not steady-state) ($ss = 0$). (see Section 2.1.2)
Fitness Assignment	fa	For fitness assignment in the evolutionary algorithm, Pareto ranking was used. (see Section 2.3.3)
Algorithm Selection	sel	A binary ($k = 2$) tournament selection was applied. (see Section 2.4.4)
Convergence Prevention	cp	No additional means for convergence prevention were used, i. e., $cp = 0$. (see Section 2.4.8)
Number of Training Cases	tc	The number of training cases used for evaluating the objective functions were $tc = 22$, where each run is granted 28 cycles in the static and 300 cycles in the dynamic case.
ct	ct	The training cases were fixed, i. e., $ct = 0$.

Table 23.2: The settings of the Aggregation-Genetic Programming experiments.

In the simulations, 16 virtual machines were running, each holding a state vector $\mathbf{s}$ with five elements. From all experiments, those with a tiered prevalence comparison performed best and, hence, will be discussed in this section. Tiered prevalence comparison is similar to a Pareto optimization which is performed level-wise. When comparing two solution candidates x_1 and x_2 , initially, the objective values of the first objective function f_1 are considered only. If one of the solution candidates has here a better value than the other, it wins. If both values are equal, we compare the second objective values in the same way, and so on. The comparator function defined in Equation 23.22 gives correctness (f_1) precedence to protocol size (f_2). Its result indicates which of the two individuals won – a negative number denotes the victory of x_1 , a positive one that x_2 is better. The tiered structure of $\text{cmp}_{F,agg}$ leads to optimal sets with few members that most often (but not always) have equal objective values and only differ in their phenotypes.

$$\text{cmp}_{F,agg}(x_1, x_2) = \begin{cases} -1 & \text{if } (f_1(x_1) < f_1(x_2)) \vee \\ & ((f_1(x_1) = f_1(x_2)) \wedge (f_2(x_1) < f_2(x_2))) \vee \\ & ((f_1(x_1) = f_1(x_2)) \wedge (f_2(x_1) = f_2(x_2)) \wedge \\ & \quad (f_3(x_1) < f_3(x_2))) \\ 1 & \text{if } (f_1(x_2) < f_1(x_1)) \vee \\ & ((f_1(x_2) = f_1(x_1)) \wedge (f_2(x_2) < f_2(x_1))) \vee \\ & ((f_1(x_2) = f_1(x_1)) \wedge (f_2(x_2) = f_2(x_1)) \wedge \\ & \quad (f_3(x_2) < f_3(x_1))) \\ 0 & \text{otherwise} \end{cases} \quad (23.22)$$

We do not need more than five memory cells in our experiments. The message size was normally one or two in all test series and if it was larger, it converged quickly to a minimum. The objective function f_3 that minimizes it thus shows no interesting behavior. It can be assumed that it will have equal characteristics like f_2 in larger problems.

Average – static

With this configuration, protocols for simple aggregates like minimum, maximum, and average can be obtained in just a few generation steps. We have used the distributed average protocol which computes $\alpha_{avg} = \bar{x}$ in many of the previous examples, for instance, in Section 23.2.2 on page 368, Section 23.2.2 on page 377, and in Fig. 23.18.1. The evolution of a static version such an algorithm is illustrated in Figure 23.19. The graphic shows how




Fig. 23.19: The evolutionary progress of the static *average* protocol.

the objective values of the first objective function (the mean square error sum) improve with the generations in twelve independent runs of the evolutionary algorithm. All runs did converge to the optimal solution previously discussed, most of them very quickly in less than 50 generations.

Figure 23.20 reveals the inter-relation between the first and second objective function for two randomly picked runs. Most often, when the accurateness of the (best known) protocols increases, the number of formula expressions also rises. These peaks in f_2 are always followed by a recession caused by stepwise improvement of the protocol efficiency by eliminating unnecessary expressions. This phenomenon is rooted in the tiered comparison that we chose: A larger but more precise protocol will always beat a smaller, less accurate one. If two protocols have equal precision, the smaller one will prevail.

Root-Of-Average – static

In our past research, we used the evolution of the *root-of-average* protocol as benchmark problem [863]. Here, a distributed average protocol for the aggregate function α_{ra} is to be evolved:

$$\alpha_{ra}(\mathbf{x}) = \sqrt{|\bar{x}|} \quad \forall x \in \mathbf{x} \quad (23.23)$$

One result of these experiments has already been sketched in Fig. 23.18.2. Figure 23.21 is a plot of eleven independent evolution runs. It also shows a solution found after only 84 generations in the quickest experiment. The values of the first objective function f_1 , denoting the mean square error, improve so quickly in all runs at the beginning that a logarithmic scale is needed to display them properly. This contrasts with the simple average protocol evolution where the measured fitness is approximately proportional to the number of generations. The reason is the underlying aggregate function which is more complicated and thus, harder to approximate. Therefore, the initial errors are much higher and even small changes in the protocols can lead to large gains in accurateness.




Fig. 23.20: The relation of f_1 and f_2 in the static *average* protocol.




Fig. 23.21: The evolutionary progress and one grown solution of the static *root-of-average* protocol.

The example solution contains a useless initialization sequence. In the experiments, it paradoxically did not vanish during the later course of the evolution although the secondary (non-functional) objective function f_2 puts pressure into the direction of smaller protocols. For the inter-relation between the first and second objective function, the same observations can be made than in the average protocol. Improvements in f_1 often cause an increase in f_2 which is followed by an almost immediate decrease, as pictured in Figure 23.22 for the 84-generation solution.




Fig. 23.22: The relation of f_1 and f_2 in the static *root-of-average* protocol.

Average – dynamic

After verifying our approach for conventional aggregation protocols with static input data, it is time to test it with dynamically changing inputs. This may turn out be a useful application and is more interesting, since creating protocols for this scenario by hand is more complicated.

So we first repeat the “average” experiment for two different scenarios with volatile input data. The first one is depicted with solid lines in Figure 23.23. Here, the *true* values of the aggregate $\alpha(\mathbf{x}(t))$ can vary in each protocol step by 1% and in one simulation by 50% in total. In the second scenario, denoted by dashed lines, these volatility measures are increased to 3% and 70% respectively. The different settings have a clear impact on the results of the error functions – the more unsteady the protocol inputs, the higher will f_1 be, as Figure 23.23 clearly illustrates. The evolved solution exhibits very simple behavior: In each protocol step, a node first computes the average of its currently known value and the new sensor input. Then, it sets the new estimate to the average of this value and the value received from its partner node. Each node sends its current sensor value. This robust basic scheme seems to work fine in a volatile environment. The course of the evolutionary process itself has not changed significantly. Also the interactions between of f_1 and f_2 stay the same, as shown in Figure 23.24.

Root-Of-Average – dynamic

Now we repeat the second experiment, evolving a protocol that computes the square root of the average, with dynamic input data. Here we follow the same approach as for the dynamic average protocol: Tests are run with the same two volatility settings as in Section 23.2.2.




Fig. 23.23: The evolutionary progress of the dynamic *average* protocol.




Fig. 23.24: The relation of f_1 and f_2 in the dynamic *average* protocol.

Figure 23.25 shows how f_1 changes by time. Like in Figure 23.21, we have to use a logarithmic scaling for f_1 to illustrate it properly. For the tests with the slower changing data (solid lines), an intermediate solution is included because the final results were too complicated to be sketched here. The evolutions with the highly dynamic input dataset however did not yield functional aggregation protocols. From this we can follow that there is a threshold of volatility from which on Genetic Programming is no longer able to breed stable formulas.




Fig. 23.25: The evolutionary progress and one grown solution of the dynamic *root-of-average* protocol.

The relation of f_1 and f_2 , outlined in Figure 23.26, complies with our expectations. In every experiment run, increasing f_1 is usually coupled to a deterioration of f_2 which means that the protocol formulas become larger and include more sub-expressions. This is followed by a recreation span where the formulas are reduced in size. After a phase of rest, where the new protocol supposable spreads throughout the population, the cycle starts all over again until the end of the evolution.

Conclusions

In this chapter, we have illustrated how Genetic Programming can be utilized for the automated synthesis of aggregation protocols. The transition to the evolution of protocols for dynamically changing input data is a step towards a new direction. Especially in applications like large-scale sensor networks, it is very hard for a software engineer to decide which protocol configuration is best. With our evolutionary approach, different solutions could be evolved for different volatility settings which can then be selected by the network according to the current situation.




Fig. 23.26: The relation of f_1 and f_2 in the dynamic *root-of-average* protocol.

Sigoa – Implementation in Java

Introduction

Today, there exist many different optimization frameworks. Some of them are dedicated to special purposes like spacecraft design [455] or trading systems [705]. Others provide multi-purpose functionality like GALib [1689], Evolutionary Objects (EO) [1690, 1691] or the Java evolutionary computation library (ECJ) [1012].

In this part of the book, we want to introduce a new approach in global optimization software, called Sigoa, the *Simple Interface for Global Optimization Algorithms*¹. Based on this library, we want to demonstrate how the optimization algorithms discussed in the previous chapters can be implemented.

We decided to use Java [1692, 1693, 1694] as programming language and runtime system for this project since it is a very common, object-oriented, and platform independent. You can find more information on Java technology either directly at the website <http://java.sun.com/> [accessed 2007-07-03] or in books like

- *Javabuch* by Krüger [1695], the best German Java learning resource in my opinion, is online available for download at <http://www.javabuch.de/> [accessed 2007-07-03].
- For the English reader, *Thinking in Java* by Eckel [1696] would be more appropriate – its free, third edition is online available at <http://www.mindview.net/Books/TIJ/> [accessed 2007-07-03].
- As same as interesting are the O'Reilly books *Java in a Nutshell* [1697] and *Java Examples in a nutshell* [1698] by Flanagan, and *Learning Java* by Niemeyer and Knudsen [1699].
- *Java ist auch eine Insel* by [1700] – another good resource written in German, is also online available at <http://www.galileocomputing.de/openbook/javainssel6/> [accessed 2007-07-03].

The source code of the binaries and the source files of the software described in this book part can be found online at <http://www.sigoa.org/>. It is not only open-source, but licensed very liberally under the LGPL (see appendix Chapter C on page 533) which allows for the integration of Sigoa into all kinds of systems, from educational to commercial, without any restrictions. Sigoa has features that aim to support optimizing complicated types of individuals which require time-consuming simulation and evaluation procedures.

Genetic Programming of real distributed algorithms is one example for such problems. In order to determine such an algorithm's fitness, you need to simulate the algorithm². The evolution progresses step by step, so at first, we will not have any algorithm that works properly for a specified problem. Some of the solution candidates whatsoever will be able to perform some of the *sub-tasks* of the problem, or will maybe solve it partly. Since they may work on some of the inputs while failing to process other inputs correctly, a single simulation run will not be sufficient. We rather execute the algorithms multiple times and then use the

¹  <http://www.sigoa.org/>

² See Section 4.10 on page 186 for further discussions.

minimum, median, average, or maximum objective values encountered. In the case of growing distributed algorithms, it is again not sufficient to simulate one processor. Instead, we need to simulate a network of many processors in order to determine the objective values³. Hence, it is simple to imagine that such a process may take some time. There are many other examples of optimization problems that involve complicated and time-consuming processes or simulations.

A framework capable of partially dealing with such aspects in an elegant manner has already been developed by the author in the past, see Weise [860], Weise and Geihs [861; 862]. With the Sigoa approach, we use our former experiences to create a software package that has a higher performance and is way more versatile: One of the key features of Sigoa is the separation of specification from implementation, which allows heavyweight implementations as required for the evolution of the distributed algorithms as well as creating lightweight optimization algorithms which do not need simulations at all – like numerical minimization or such and such. This clear division not only allows for implementing all the optimization algorithms introduced in the previous parts but is good basis for including new, experimental methods that may have not been discussed yet.

Before starting with the specification of the Sigoa approach, we performed a detailed study on different global optimization methods and evolutionary algorithms. Additionally, we used the lessons learned from designing the Dgpf system to write down the following major requirements:

24.1 Requirements Analysis

24.1.1 Multi-Objectivity

Almost all real-world problems involve contradicting objectives. A distributed algorithm evolved should, for example, be efficient and yet simple. It should consume not much memory and involve as little as possible communication between different processors but on the other hand should ensure proper functionality and be robust against lost or erroneous messages. The first requirement of an optimization framework is thus multi-objectivity.

24.1.2 Separation of Specification and Implementation

It should easily be possible to adapt the optimization framework to other problems or problem domains. The ability to replace the solution candidate representation forms is therefore necessary. Furthermore, the API must allow the implementation of all optimization algorithms discussed in the previous chapters in an easy and elegant manner. It should further be modular, since most of the optimization algorithms also consist of different sub-algorithms, as we have seen for example in Chapter 2 on page 67.

From this requirement we deduce that the software architecture used for the whole framework should be component based. Each component should communicate with the others only through clearly specified interfaces. This way, each module will be exchangeable and may be even represented by proxies or such and such, granting a maximum of extensibility. If we define a general interface for selection, we could modify the SPEA-algorithm (see ?? on page ??) which originally uses tournament selection to use another selection algorithm.

Hence, we will define Java-interfaces for all parts of optimization algorithms such as fitness assignment or clustering methods used for pruning the optimal sets. By doing so, we reach a separation of the specification from the implementation. For all interfaces we will provide a reference implementation which can easily be exchanged, allowing for different levels of complexity in the realizations.

³ In Section 23.2.2 on page 366 you can find good example for this issue.

24.1.3 Separation of Concerns

An optimization system consists not only of the optimization algorithms themselves. It needs interfaces to simulators. If it is distributed, there must be a communication subsystem. Even if the optimization system is not distributed, we will most likely make use of parallelism since the processors inside of nowadays off-the-shelf PCs already offer supportive hyper-threading or dual-core technology [1701, 1472]. If Sigoa is utilized by multiple other software systems which transfer optimization tasks to it, security issues arise. These aspects are orthogonal to the mathematical task of optimizing and should therefore be specified at different places and clearly be separated from the pure algorithms. Best practice commands to already consider such aspects in the earliest software design phase of every project and thus, also in the Sigoa library.

24.1.4 Support for Pluggable Simulations and Introspection

In most real-world scenarios, simulations are needed to evaluate the objective values of the solution candidates. If we use the framework for multiple problem domains, we will need to exchange these simulations or even want to rely on external modules. In some cases, the value of an objective function is an aggregate of everything what happened during the simulation. Therefore, they need a clear insight into what is going. Since we separate the objective functions from the simulations by clearly specified interfaces (as discussed in Section 24.1.3), these interfaces need to provide this required functionality of introspection.

In the use case of evolving a distributed algorithm, we can visualize the combination with the separation of concerns and introspective simulations: Besides working correctly, a distributed algorithm should use as few messages as possible or at least has stable demands considering the bandwidth on the communication channel. We therefore could write an objective function which inspects the number of messages underway in the simulation and computes a long-term average and variance. The simulation itself then does not need to be aware of that; it simply has to offer the functionality of counting the messages currently in transmission. The catch is that we can now replace the objective function by another one that maybe puts the pressure a little bit differently, leading to better results, without modifying the simulation. On the other hand, we can also use different simulation models – for example one where transmission errors can occur and one where this is not the case – without touching the objective function.

24.1.5 Distribution utilities

As already said, there are many applications where the simulations are very complicated and therefore, our architecture should allow us to distribute the arising workload to a network of many computers. The optimization process then can run significantly faster because many optimization techniques (especially evolutionary algorithms) are very suitable for parallel and distributed execution as discussed in Chapter 17 on page 255.

24.2 Architecture

We want to design the Sigoa optimization system based on these requirements. In this book part, we have assigned different chapters to the classes of different components of Sigoa and their sub-algorithms. By specifying interfaces for all aspects of optimization and implementing them elsewhere, the versatility to exchange all components is granted, so customized optimizers can be built to obtain the best results for different problem domains. Furthermore, interfaces allow us to implement components in different levels of detail: there may be applications where the evaluation of objective functions involves massive simulations

(like genetic programming) and applications, where the simple evaluation of mathematical functions enough (like numerical minimizing). In the latter case, using a system that provides extended support for simulations may result in performance degeneration since a lot of useless work is performed. If the mechanism that computes the objective values could be exchanged, an optimal approach can be used in each case.

As result from these considerations, we divide the Sigoa architecture in `org.sigoa` into two main packages: `org.sigoa.spec` contains the specifications and `org.sigoa.refimpl` a reference implementation. Figure 24.1 illustrates this top-level package hierarchy.




Fig. 24.1: The top-level packages of the Sigoa optimization system.

The specification of the functionality of Sigoa is given by interfaces and a few basic utility classes only. It is independent from any library or other software system and does not require prerequisites. The interfaces can therefore also be implemented as wrappers that bridge to other, existing optimizing systems. Most of these specification interfaces inherit from `java.io.Serializable` and hence can be serialized using the Java Object Serialization mechanisms⁴. This way, we provide the foundation for creating snapshots of a running optimization algorithm which would allow for starting, stopping, restarting, and migrating them.

The reference implementation uses an additional software package called Sfc, the *Java Software Foundation Classes* – a LGPL-licensed open-source library available under the same URL as Sigoa that provides some useful classes for tasks that are needed in many applications like enhanced IO, XML support, extended and more efficient implementations of the Java Collection Framework⁵-interfaces and so on. This utility collection is not directly linked to optimization algorithms but provides valuable services that ease the implementation of the Sigoa components.

The package hierarchy of the reference implementation is identical to the one of the specifications. The package `org.sigoa.spec.gp.reproduction`, for example, contains the definition of mutation and crossover operations whereas the package `org.sigoa.refimpl.gp.reproduction` contains the reference implementation of these operations.




Fig. 24.2: The subsystem specification of the optimization framework.

24.3 Subsystems

As illustrated in Figure 24.2, the Sigoa framework is constituted by nine subsystems:

1. The **adaptation** package contains mechanisms that allow components to adapt themselves to a given situation based on rules. This can be used for example by optimization algorithms in order to adjust their parameters. A very simple application is the termination criterion⁶: a rule could be defined that terminates an algorithm after a given time.
2. In the **clustering**-package, interfaces for clustering-algorithms (as defined in Chapter 28 on page 467) are specified. Clustering algorithms can be used to reduce a large set of solution candidates to a smaller one without loss of generality.
3. One way for optimization algorithms to report their status and statistics to the outside world would be via events. As already said, we do not treat the optimization process as a mere mathematical procedure – it will always be part of some application. As such, not only the final results are interesting but also status messages and statistic evaluations of its progress. The **events** package defines interfaces for events that can be generated and may contain such information.
4. The largest subsystem is the **go** package, where all components and sub-algorithms for global optimization are specified. Here you can find the interface specifications that cover the all the algorithmic and mathematical functionality of global optimization algorithms.
5. In the **jobssystem** package, we place the specification of the means to *run* optimization algorithms. An optimizer may be parallelized or run sequentially and therefore may use multiple threads. The algorithm itself should be separated from the parallelism issues. Applying the definitions of the **jobssystem** package, optimizers may divide their work into

⁴ <http://java.sun.com/javase/6/docs/technotes/guides/serialization/> [accessed 2007-07-03]

⁵ <http://java.sun.com/javase/6/docs/technotes/guides/collections/> [accessed 2007-07-03]

⁶ see Section 1.3.4 on page 35

parallelizable pieces which can be executed as *jobs*. Such jobs are then handled by the job system, which decides if they should be run in different threads or performed sequentially. This way, it is also possible to manage multiple optimization algorithms in parallel and to specify which one will be assigned to how many processors. The implementations of the job system specifications could also perform accounting and performance monitoring of the work load.

6. The concept of pipes defined in the `pipe` package is a very mighty approach for realizing optimization. It does not only allow separating the different components of an optimizer completely – totally new components, like statistic monitors can also easily be inserted into a system with minimum changes.
7. The job system enables Sigoa to handle multiple optimization requests at once. Since it is a plain component interface, these requests may come from anywhere, maybe even from a web service interface built on top of it. It must somehow be ensured that such requests do not interfere or even perform harmful or otherwise malicious actions. Therefore, a security concept is mandatory. In the `security` package we specify simple interfaces that build on the Java Security Technology⁷.
8. The behavior of solution candidates is often simulated in order to determine their objective values. The `simulation` package provides interfaces that specify how simulations can be accessed, made available, and are managed.
9. Stochastic evaluations are a useful tool for optimization algorithms. As already said, the application using the Sigoa system may regularly need information about the progress, which normally can only be given in form of some sort of statistical evaluation. This data may also be used by the optimization algorithms themselves or by adaptation rules. Furthermore, most the global optimization methods discussed here are randomized algorithms. They thus need random number generators as introduced in Section 27.7 on page 459.

⁷ <http://java.sun.com/javase/6/docs/technotes/guides/security> [accessed 2007-07-03]

Examples

But before we are going into detail about the different packages and utilities of the Sigoa software, we will present some application examples. These give a straightforward insight into the usage and customization of the Sigoa components which most probably is good enough to apply them to other problems. A more specific discussion of the Sigoa packages following after this chapter then rounds up the view on this novel optimization system.

25.1 The 2007 DATA-MINING-CUP

As an application example for genetic algorithms, the 2007 DATA-MINING-CUP Contest, has been introduced in Section 21.1.2 on page 318. We strongly recommend reading this section first. We there have discussed the basic principles behind the challenge and the structure of one possible solution to it. Here we will only show how these ideas can be realized easily using the Sigoa framework.

The objective of the contest is to classify a set of 50 000 data vectors containing 20 features (from which only 17 are relevant) each into one of the three groups **A**, **B**, and **N**. In order to build classifiers that do so, another 50 000 datasets with already known classifications are available as training data. Thus, let us start by representing the three possible classification results using a simple Java `enum` like in Listing 25.1.

Our approach in Section 21.1.2 was to solve the task using a modified version of Learning Classifier Systems C . For the contest, a function $P(C)$ denoting the profit that could be gained with a classifier C was already defined (see Equation 21.1). Thus, we simply strip the LCSs from their *learning* capability and directly maximize the profit directly.

25.1.1 The Phenotype

The problem space $\mathbb{X}$ was thus composed of mere classifier systems, the phenotypes of a genetic algorithm. They consist of a set of rules `m_rules` (the single classifiers) which we can

```
1 public enum EClasses {
2     /** class A */
3     A,
4     /** class B */
5     B,
6     /** class N */
7     N;
8 }
```

Listing 25.1: The enum `EClasses` with the possible DMC 2007 classifications.

represent as `byte` arrays containing the conditions and rule outputs `m_results` (instances of `EClasses`).

Listing 25.2 illustrates how the method `classify` works which classifies a set of 17 relevant features (stored in the array `data`) into one of the three possible `EClasses` instances. It iterates over all the rules in `m_rules` and checks if rule `m_rules[i]` fits perfectly according to the definitions in Table 21.2 on page 322. If so, the corresponding classification `m_results[i]` is returned. `classify` keeps also track on the rule which has the fewest violated conditions in the variables `lec` and `leci`. If no perfectly matching rule could be found, the $\frac{1}{5}$ -threshold mentioned in Section 21.1.2 is checked: if `lec` $\leq$ 3, the classification `m_results[leci]` belonging to the rule with the least violations is returned. Otherwise, the class `N` represented by `EClasses.N` is assigned to the data sample.

So this is basically what a *phenotype* can look like in Sigoa – you can clearly see that, except from implementing `java.io.Serializable`, no further requirements are imposed on its structure. The method `classify` is not mandatory, it is will just be part of the evaluation in this particular optimization problem; other problems may need totally other functionality.

25.1.2 The Genotype and the Embryogeny

The *genotype* that belongs to the phenotypic individual representations is a variable-length a bit string. Such genomes have been discussed in Section 3.5 on page 121 extensively. In Figure 21.4, we have introduced the genotype-phenotype mapping in this particular application: since there are four possible conditions and 17 conditions plus three possible classifications (`A`, `B`, and `N`) per rule, we need $17 * 2 + 2 = 36$ bits to encode a single classifier which is the *granularity*, i. e., the gene size of our genome. A classifier system in turn may consist of an arbitrary number of such classifiers.

In Sigoa, we can represent variable-length as well as fixed-length bit strings as byte arrays (`byte[]`) for which predefined creation, mutation, and crossover operators exist. Therefore, we do not have to deal with the reproduction operations directly and can concentrate on the translation of a genotype $g \in \text{byte}[]$ into a corresponding phenotype which is an instance of `ClassifierSystem`. Such translations is called genotype-phenotype mapping (see Section 3.6 on page 122) or artificial embryogeny (discussed in Section 3.6) for which Sigoa offers a core interface `IEmbryogeny` (see ?? on page ??) and a reference implementation `Embryogeny` (see ?? on page ??) along with an extension for bitstrings, `BitStringEmbryogeny` (see ?? on page ??) which provides special streams for input and output of structured data from and to bit strings. We simply need to extend this class by providing (at least) the transformation function `gpm` from genotypes to phenotypes and (optionally) vice versa. Listing 25.3 shows this extension in form of the class `ClassifierEmbryogeny`. The constant `CLASSIFIER_EMBRYOGENY` provides a globally shared instance of this new embryogeny.

25.1.3 The Simulation

So now we need to find out how an evolved classifier system C behaves. Therefore we can use the provided test datasets or better, only a good share of them while saving the rest in order to check if our classifier system generalize well. For these training sets, we built a matrix $M(C)$ where the columns denote the classification results delivered by C and the rows contain the true classes. For determining the zero-based indices into this matrix we use the method `ordinal()` of the `EClasses` enum, i. e., $m_{2,1}$ would represent those elements in class `N` that were miss-classified into group `B` – 2799 in the example matrix M_{ex} of Equation 25.1. From M_{ex} , we can furthermore read that 1087 of the samples belonging to class `B` were correctly classified whereas 777 were assigned to class `A` and 1462 to class `N`.

$$M_{ex}(C) = \begin{pmatrix} 4062 & 856 & 3794 \\ 777 & 1087 & 1462 \\ 5484 & 2799 & 29\,679 \end{pmatrix} \quad (25.1)$$

```

1 public class ClassifierSystem extends JavaTexttable implements
   Serializable {
2     ...
3     private final byte[][] m_rules;
4     private final EClasses[] m_results;
5     ...
6     public ClassifierSystem(final byte[][] rules, final EClasses[]
       results) {
7         super();
8         this.m_results = results;
9         this.m_rules   = rules; }
10    ...
11    public EClasses classify(final byte[] data) {
12        byte[][] rules;
13        byte[] rule;
14        int i, j, ec, lec, leci;
15
16        rules = this.m_rules;
17        lec   = Integer.MAX_VALUE;
18        leci  = 0;
19
20        main: for (i = (rules.length - 1); i >= 0; i--) {
21            rule = rules[i];
22            ec = 0;
23            for (j = (rule.length - 1); j >= 0; j--) {
24                switch (rule[j]) {
25                    case 0: {
26                        if (data[j] != 0)
27                            if ((++ec) > 3) continue main;
28                        break;
29                    }
30                    case 1: {
31                        if (data[j] < 1) // != 1
32                            if ((++ec) > 3) continue main;
33                        break;
34                    }
35                    case 2: {
36                        if (data[j] <= 1) // <= 0)
37                            if ((++ec) > 3) continue main;
38                        break;
39                    }
40                }
41            }
42            if (ec <= 0) return this.m_results[i];
43            if (ec < lec) {
44                lec = ec;
45                leci = i;
46            }
47        }
48        if (lec <= 3) return this.m_results[leci];
49        return EClasses.N;
50    }
51    ...
52 }

```

Listing 25.2: The structure of our DMC 2007 classifier system.

```

1 public class ClassifierEmbryogeny extends
   BitStringEmbryogeny<ClassifierSystem> {
2     /** the classes */
3     private static final EClasses[] CLASSES = EClasses.values();
4     /** The globally shared instance */
5     public static final IEmbryogeny<byte[], ClassifierSystem>
        CLASSIFIER_EMBRYOGENY = new ClassifierEmbryogeny();
6     ..
7     /** This method is supposed to compute an instance of
8      * the phenotype from an instance of the genotype.
9      * @param genotype The genotype instance to breed a
10     * phenotype from.
11     * @return The phenotype hatched from the genotype. */
12    @Override
13    public ClassifierSystem hatch(final byte[] genotype) {
14        int            i, j, c;
15        byte[][]       rules;
16        byte[]         rule;
17        EClasses[]     results;
18        BitStringInputStream bis;
19
20        if (genotype == null)
21            throw new NullPointerException();
22        c      = ((genotype.length * 8) / 36);
23        rules  = new byte[c][17];
24        results = new EClasses[c];
25        bis    = this.acquireBitStringInputStream();
26        bis.init(genotype);
27
28        for (i = 0; i < c; i++) {
29            rule = rules[i];
30            for (j = 0; j < 17; j++) {
31                rule[j] = (byte) (bis.readBits(2));
32            }
33            results[i] = (CLASSES[bis.readBits(2) % 3]);
34        }
35        this.releaseBitStringInputStream(bis);
36        return new ClassifierSystem(rules, results);
37    }
38    ...
39 }

```

Listing 25.3: The embryogeny component of our DMC 2007 contribution.

From such matrices, we can easily compute the profit function $P(C)$ as well as other features, like how many **As**, **Bs**, and **Ns** were classified incorrectly. What we basically do here is to simulate the behavior of the classifiers. And for simulations, Sigoa provides the interface `ISimulation` (see ?? on page ??) and its standard implementation `Simulation` (see ?? on page ??). This default implementation just needs to be extended so it uses the training samples, which we load somewhere else (in a class called `Datasets`), and computes M . Therefore, overriding the method `beginIndividual` is sufficient and other changes are not needed.

Listing 25.4 shows the most important code of the new class `ClassificationSimulation`. In order to allow us to publish the new simulation in the simulation manager of the optimization job system, we also provide a globally shared factory in form of an `ISimulationProvider`-instance with the constant `PROVIDER` in line 3.

```

1 public class ClassificationSimulation extends
    Simulation<ClassifierSystem> {
2     /** the shared provider for this simulation */
3     public static final ISimulationProvider PROVIDER = new
        SimulationProvider(ClassificationSimulation.class);
4     /** the matrix M(C) */
5     private final int[][] m_classifications;
6     ...
7     public ClassificationSimulation() {
8         super();
9         this.m_classifications = new int[3][3]; }
10
11     /** Here the matrix M(C) is computed
12     * @param what The classifier C to be simulated. */
13     @Override
14     public void beginIndividual(final ClassifierSystem what)
15     {
16         int i;
17         int[][] x = this.m_classifications;
18         super.beginIndividual(what);
19         for (i = (x.length - 1); i >= 0; i--)
20             Arrays.fill(x[i], 0);
21         for (i = (DATA.length - 1); i >= 0; i--)
22             x[CLASSES[i].ordinal()][ what.classify(DATA[i]).ordinal() ]++;
23     }
24     ...
25     /** Compute the profit value P(C). */
26     public int getProfit() {
27         int[][] data = this.m_classifications;
28         return (3 * data[0][0]) + (6 * data[1][1]) - (data[2][0] +
            data[2][1] + data[0][1] + data[1][0]);
29     }
30 }

```

Listing 25.4: The simulation for testing the DMC 2007 classifier system.

25.1.4 The Objective Functions

On the foundation of the new simulation for classifier system, we can define the objective functions that should guide the evolution. In Section 21.1.2 on page 323 we have introduced the two most important objective functions: one that minimizes $f_1(C) = -P(C)$ and hence, maximizes the profit, and $f_2(C) = |C|$ which minimizes the number of rules in the classifier system.

All objective functions in Sigoa are instances of the interface `IObjectiveFunction` (see ?? on page ??). They can be derived from its default implementation `ObjectiveFunction` (see ?? on page ??) which implements the basic functionality so only the real mathematical computations need to be added.

In Listing 25.5, we implement f_1 . Therefore, the method `endEvaluation` needs to be overridden. Here we store negated profit into a state record which is used by the optimization system to compute the objective value and to store it into the individual records later on.

The only remaining question is: How will the optimizer system know that our objective function needs an instance of `ClassificationSimulation` and that it has to call its method `beginIndividual` beforehand? The answer is relatively simple: In line 3, we have defined an instance of `SimulationProvider` for the `ClassificationSimulation`. This provider will later be introduced to the optimization job system. It uses `ClassificationSimulation.class` as identifier per default. With the method `getRequiredSimulationId` on line 16, we tell the

```

1 public class ProfitObjectiveFunction extends
    ObjectiveFunction<ClassifierSystem, ObjectiveState, Serializable,
    ClassificationSimulation> {
2
3     /** This method is called after any simulation/
4     * evaluation is performed. It stores the negated
5     * profit  $-P(C)$  into the state-variable - that's all.*/
6     @Override
7     public void endEvaluation(final ClassifierSystem individual, final
        ObjectiveState state, final Serializable staticState, final
        ClassificationSimulation simulation) {
8         state.setObjectiveValue(-simulation.getProfit());
9     }
10
11     /**
12     * Obtain the ID of the required simulator.
13     * @return The ID=class of our simulator */
14     @Override
15     public Serializable getRequiredSimulationId() {
16         return ClassificationSimulation.class;
17     }
18 }

```

Listing 25.5: The profit objective function $f_1(C) = -P(C)$ for the DMC 2007.

```

1 public class SizeObjectiveFunction extends
    ObjectiveFunction<ClassifierSystem, ObjectiveState, Serializable,
    ISimulation<ClassifierSystem>> {
2
3     /** This method is called after any simulation/
4     * evaluation is performed. It stores the size of
5     * the \ClassS\  $|C|$  into the state-
6     * variable - that's all. */
7     @Override
8     public void endEvaluation(final ClassifierSystem individual, final
        ObjectiveState state, final Serializable staticState, final
        ISimulation<ClassifierSystem> simulation) {
9         state.setObjectiveValue(Math.max(individual.getSize(), 3));
10    }

```

Listing 25.6: The size objective function $f_2(C) = |C|$ for the DMC 2007.

job system that we need a simulation which is made available by a provider with exactly this ID. Before passing the simulation to our objective function, the job system will call its `beginIndividual` method which, in turn, builds the matrix $M(C)$ holding the information needed for its `getProfit` method. Now we can query the profit from this simulation.

For the secondary objective function f_2 defined in Listing 25.6, we do not need any simulation. Instead, we directly query the number of rules in the classifier system via the method `getSize`. In Listing 25.2, we have omitted this routine for space reasons, it simply returns `m_rules.length`. Again, this value is stored into the state record passed in from the evaluator component of the job system which will then do the rest of the work.

25.1.5 The Evolution Process

Now the work is almost done, we just need to start the optimization process. Listing 25.7 presents a `main`-method which is called on startup of a Java program and does so. Therefore, we first have to decide which global optimization algorithm should be used and pick `ElitsitEA`¹, an elitist evolutionary algorithm (per default steady-state) with a population size of $10 * 1024$ and mutation and crossover rates of 0.4 in line 8.

Then we construct an `IOptimizationInfo`-record with all the information that will guide the evolution². Part of this information is how the solution candidates should be evaluated. For this, we use an instance of `Evaluator`³ (line 15) which is equipped with a `List` containing the two new objective functions from 10 and 12. We furthermore tell the system to perform a pure Pareto-optimization as discussed in Section 1.2.4 on page 13 by passing the globally shared instance of `ParetoComparator`⁴ (line 16) into the info record. We then define that our embryogeny component should be used to translate the bit string genotypes into `ClassifierSystem` phenotypes in line 17. These genotypes are produced by the default reproduction operators for variable-length bit string genomes⁵ added in lines 18 to 20. All of them are created with a *granularity* of 36 which means that it is ensured that all genotypes have sizes of multiples of 36 bits and that all crossover operations only occur at such 36 bit boundaries.

After this is done, we instantiate `SimulationManager`⁶ and publish the new simulation that we have developed in Section 25.1.3 on page 398 by adding its provider to the simulation manager in line 27. The job system created in line 28 allows the evaluator to access the simulation manager, an instance of the interface `ISimulationManager`⁷. The evaluator will then ask its objective functions which simulations they need – in our case a `ClassificationSimulation` – and then query the simulation manager to provide them.

In line 28, we decided to use a multi-processor job system which is capable of transparently parallelizing the optimization process. The different types of job systems which are instances of the interface `IJobSystem` specified in ?? on page ?? are discussed in ?? on page ?. We add our optimizer to the system in line 29 and finally start it in 30. Since we have added no termination criterion, the system will basically run forever in this example.

In order to get information on its progress, we have provided two special output pipes (see ?? on page ??) in lines 23 and 24 to the optimizer's non-prevalence pipe. Through this pipe, in each generation all non-prevailed (in our case, non-dominated) individuals will be written and thus, pass our two pipes. In each generation, new text files with information about them are created. The first one, the `IndividualPrinterPipe`, uses the directory *c* and creates files that start with a *c* followed by the current generation index. It writes down the full information about all individuals. From these files, we can later easily reconstruct the complete individuals and, for instance, integrate them into the real applications. The second printer pipe, an instance of `ObjectivePrinterPipe`, stores only the objective values in a comma-separated-values format. The output files are put into the directories *bo* and also start with *bo* followed by the generation index. Such files are especially useful for getting a quick overview on how the evolution progresses. Later, they may also read into spread sheets or graphical tools in order to produce fancy diagrams like Figure 21.5 on page 324.

¹ see ?? on page ??

² `IOptimizationInfo` is discussed in ?? on page ??, its reference implementation `OptimizationInfo` in ?? on page ??.

³ The class `Evaluator`, discussed in ?? on page ??, is the default implementation of the interface `IEvaluator` specified in ?? on page ??.

⁴ The class `ParetoComparator`, elaborated on in ?? on page ??, implements the interface `IComparator` defined in ?? on page ??.

⁵ These operations are introduced in ?? on page ?? implement the interfaces `ICreator`, `IMutator`, and `ICrossover` specified in ?? on page ??.

⁶ see ?? on page ??

⁷ see ?? on page ??

```

1 public static void main(String[] args) {
2     EA<byte[], ClassifierSystem> opt;
3     IOptimizationInfo<byte[], ClassifierSystem> oi;
4     IJobSystem s;
5     SimulationManager m;
6     List<IOjectiveFunction<ClassifierSystem, ?, ?,
7         ISimulation<ClassifierSystem>>> l;
8
9     opt = new ElitistEA<byte[], ClassifierSystem>(10 * 1024, 0.4d,
10        0.4d);
11
12     l = new ArrayList<IOjectiveFunction<ClassifierSystem, ?, ?,
13        ISimulation<ClassifierSystem>>>();
14     l.add(new ProfitObjectiveFunction());
15     l.add(new SizeObjectiveFunction());
16
17     oi = new OptimizationInfo<byte[], ClassifierSystem>(
18        new Evaluator<ClassifierSystem>(l),
19        ParetoComparator.PARETO_COMPARATOR,
20        ClassifierEmbryogeny.CLASSIFIER_EMBRYOGENY,
21        new VariableLengthBitStringCreator(36),
22        new VariableLengthBitStringMutator(36),
23        new VariableLengthBitStringNPointCrossover(36),
24        null);
25
26     opt.addNonPrevailedPipe(new IndividualPrinterPipe<byte[],
27        ClassifierSystem>(new FileTextWriterProvider(new
28        File("c"), "c"), false));
29     opt.addNonPrevailedPipe(new ObjectivePrinterPipe<byte[],
30        ClassifierSystem>(new FileTextWriterProvider(new File("bo"),
31        "bo"), false));
32
33     m = new SimulationManager();
34     m.addProvider(ClassificationSimulation.PROVIDER);
35     s = new MultiProcessorJobSystem(m);
36     s.executeOptimization(opt, new JobInfo<byte[],
37        ClassifierSystem>(oi));
38     s.start();
39 }

```

Listing 25.7: A main method that runs the evolution for the 2007 DMC.

Part IV

Background

Set Theory

Set theory¹ [1702, 1703, 1704] is an important part of the mathematical theory. Numerous other disciplines like algebra, analysis and topology are based up on it. Since set theory (and the topics to follow) is not the topic of this book but a mere prerequisite, this chapter (and the ones to follow) will just briefly introduce it. We make heavy use of wild definitions and in some cases even use roughly cut stuff short. More information on the topics discussed can be retrieved from the literature references.

Set theory can be divided into naïve set theory² and axiomatic set theory³. The first approach, the naïve set theory, is inconsistent and therefore not regarded in this book.

Definition 26.1 (Set). A set is a collection of objects considered as a whole⁴. The objects of a set are called elements or members. They can be anything, from numbers and vectors, to complex data structures, algorithms, or even other sets. Sets are conventionally denoted with capital letters, A , B , C , etc. while their elements are usually referred to with small letters a , b , c .

26.1 Set Membership

The expression $a \in A$ means that the element a is a member of the set A while $y \notin A$ means that y is not a member of A . A set can contain an element only once. There are three common forms to define sets:

- With their elements in braces: $A = \{1, 2, 3\}$ defines a set A containing the three elements 1, 2, and 3.
- The same set can be specified using logical operators to describe its elements: $\forall b \in \mathbb{N} : (b \geq 1) \wedge (b < 4) \Leftrightarrow b \in B$.
- A shortcut for the previous form is to denote the logical expression in braces, like $C = \{(c \geq 1) \wedge (c < 4), c \in \mathbb{N}\}$.

The cardinality of a set A is written as $|A|$ and stands for the count of elements in the set.

26.2 Relations between Sets

Two sets A and B are said to be equal, written $A = B$, if they have the same members. They are not equal ($A \neq B$) if either a member of A is not an element of B or an element

¹ http://en.wikipedia.org/wiki/Set_theory [accessed 2007-07-03]

² http://en.wikipedia.org/wiki/Naive_set_theory [accessed 2007-07-03]

³ http://en.wikipedia.org/wiki/Axiomatic_set_theory [accessed 2007-07-03]

⁴ http://en.wikipedia.org/wiki/Set_%28mathematics%29 [accessed 2007-07-03]

of B is not a member of A . If all elements of the set A are also elements of the set B , A is called subset of B and B is the superset of A . We write $A \subset B$ if A is a (true) subset of but not equal to B . $A \subseteq B$ means the A is a subset of B and may be equal to B . If A is no subset of but may be equal to B , $A \not\subseteq B$ is written. $A \not\subset B$ means that A is neither a subset of nor equal to B .

$$A = B \equiv x \in A \Leftrightarrow x \in B \quad (26.1)$$

$$A \neq B \equiv (\exists x : x \in A \wedge x \notin B) \vee (\exists y : y \in B \wedge y \notin A) \quad (26.2)$$

$$A \subseteq B \equiv x \in A \Rightarrow x \in B \quad (26.3)$$

$$A \subset B \equiv A \subseteq B \wedge \exists y : y \in B \wedge y \notin A \quad (26.4)$$

$$A \not\subseteq B \equiv \exists x : x \in A \wedge x \notin B \quad (26.5)$$

$$A \not\subset B \equiv (A = B) \vee (\exists x : x \in A \wedge x \notin B) \quad (26.6)$$

26.3 Special Sets

Special sets used in the context of this book are

- The empty set $\emptyset = \{\}$ contains no elements ($|\emptyset| = 0$).
- The natural numbers⁵ $\mathbb{N}$ include all whole numbers bigger than 0. ($\mathbb{N} = \{1, 2, 3, \dots\}$)
- The natural numbers including 0 ($\mathbb{N}_0$) include all whole numbers bigger than or equal to 0. ($\mathbb{N}_0 = \{1, 2, 3, \dots\}$)
- $\mathbb{Z}$ is the set of all integers, positive and negative. ($\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$)
- The rational numbers⁶ $\mathbb{Q}$ are defined as $\mathbb{Q} = \{\frac{a}{b} : a, b \in \mathbb{Z}, b \neq 0\}$.
- All real numbers⁷ $\mathbb{R}$ include all rational and irrational numbers (such as $\sqrt{2}$).
- $\mathbb{R}^+$ denotes the positive real numbers including 0: ($\mathbb{R}^+ = [0, \infty)$).
- $\mathbb{C}$ is the set of complex numbers⁸. When needed in the context of this book, we abbreviate the imaginary unit with i , and the real and imaginary parts of a complex number z with $\text{Re}z$ and $\text{Im}z$.

$$\mathbb{N} \subset \mathbb{N}_0 \subset \mathbb{Z} \subset \mathbb{Q} \subset \mathbb{R} \subset \mathbb{C} \quad (26.7)$$

$$\mathbb{N} \subset \mathbb{N}_0 \subset \mathbb{R}^+ \subset \mathbb{R} \subset \mathbb{C} \quad (26.8)$$

For these numerical sets, special subsets, so called intervals, can be specified. $[1, 5)$ is a set which contains all the numbers starting from (including) 1 up to (exclusively) 5. $(1, 5]$ on the other hand contains all numbers bigger than 1 and up to inclusively 5. In order to avoid ambiguities, such sets will always be used in a context where it is clear if the numbers in the set are natural or real.

26.4 Operations on Sets


Let us now define the possible unary and binary operations on sets, some of which are illustrated in Figure 26.1.

⁵ http://en.wikipedia.org/wiki/Natural_numbers [accessed 2008-01-28]

⁶ http://en.wikipedia.org/wiki/Rational_number [accessed 2008-01-28]

⁷ http://en.wikipedia.org/wiki/Real_numbers [accessed 2008-01-28]

⁸ http://en.wikipedia.org/wiki/Complex_number [accessed 2008-01-29]

Fig. 26.1: Set operations performed on sets A and B inside a set $\mathbb{A}$

Definition 26.2 (Set Union). The union⁹ C of two sets A and B is written as $A \cup B$ and contains all the objects that are element of at least one of the sets.

$$C = A \cup B \Leftrightarrow ((c \in A) \vee (c \in B) \Leftrightarrow (c \in C)) \quad (26.9)$$

$$A \cup B = B \cup A \quad (26.10)$$

$$A \cup \emptyset = A \quad (26.11)$$

$$A \cup A = A \quad (26.12)$$

$$A \subseteq A \cup B \quad (26.13)$$

Definition 26.3 (Set Intersection). The intersection¹⁰ D of two sets A and B , denoted by $A \cap B$, contains all the objects that are elements of both of the sets. If $A \cap B = \emptyset$, meaning that A and B have no elements in common, they are called *disjoint*.

$$D = A \cap B \Leftrightarrow ((d \in A) \wedge (d \in B) \Leftrightarrow (d \in D)) \quad (26.14)$$

$$A \cap B = B \cap A \quad (26.15)$$

$$A \cap \emptyset = \emptyset \quad (26.16)$$

$$A \cap A = A \quad (26.17)$$

$$A \cap B \subseteq A \quad (26.18)$$

Definition 26.4 (Set Difference). The difference E of two sets A and B , $A \setminus B$, contains the objects that are element of A but not of B .

$$E = A \setminus B \Leftrightarrow ((e \in A) \wedge (e \notin B) \Leftrightarrow (e \in E)) \quad (26.19)$$

$$A \setminus \emptyset = A \quad (26.20)$$

$$\emptyset \setminus A = \emptyset \quad (26.21)$$

$$A \setminus A = \emptyset \quad (26.22)$$

$$A \setminus B \subseteq A \quad (26.23)$$

⁹ http://en.wikipedia.org/wiki/Union_%28set_theory%29 [accessed 2007-07-03]

¹⁰ http://en.wikipedia.org/wiki/Intersection_%28set_theory%29 [accessed 2007-07-03]

Definition 26.5 (Set Complement). The complementary set $\overline{A}$ of the set A in a set $\mathbb{A}$ includes all the elements which are in $\mathbb{A}$ but not element of A :

$$A \subseteq \mathbb{A} \Rightarrow \overline{A} = \mathbb{A} \setminus A \quad (26.24)$$

Definition 26.6 (Cartesian Product). The Cartesian product¹¹ P of two sets A and B , denoted $P = A \times B$ is the set of all ordered pairs (a, b) whose first component is an element from A and the second is an element of B .

$$P = A \times B \Leftrightarrow P = \{(a, b) : a \in A, b \in B\} \quad (26.25)$$

$$A^n = \underbrace{A \times A \times \dots \times A}_{n \text{ times}} \quad (26.26)$$

Definition 26.7 (Countable Set). A set S is called countable¹² if there exists an injective function¹³ $\exists f : S \mapsto \mathbb{N}$.

Definition 26.8 (Uncountable Set). A set is uncountable if it is not countable, i. e., no such function exists for the set. $\mathbb{N}$, $\mathbb{Z}$, and $\mathbb{Q}$ are countable, $\mathbb{R}$ and $\mathbb{R}^+$ are not.

Definition 26.9 (Power Set). The power set¹⁴ $\mathcal{P}(A)$ of the set A is the set of all subsets of A .

$$\forall p \in \mathcal{P}(A) \Leftrightarrow p \subseteq A \quad (26.27)$$

26.5 Tuples

Definition 26.10 (Type). A type is a set of values that a variable, constant, function, or similar entity may take on.

We can, for instance, specify the type $T = \{1, 2, 3\}$. A variable x which is an instance of this type then can take on the values 1, 2, or 3.

Definition 26.11 (Tuple). A tuple¹⁵ is an ordered, finite sequence of elements, where each element is an instance of a certain type.

To each position in i a tuple t , a type T_i is assigned. The element $t[i]$ at a position i must then be an element of T_i . Other than sets, tuples may contain the same element twice. Since every item of a tuple may be of a different type, $(Monday, 23, \{a, b, c\})$ is a valid tuple.

In the context of this book, we define tuples with parenthesis like (a, b, c) whereas sets are specified using braces $\{a, b, c\}$.

Definition 26.12 (Tuple Type). To formalize this relation, we define the tuple type T . T specifies the basic sets for the elements of its tuples. If a tuple t meets the constraints imposed to its values by T , we can write $t \in T$ which means that the tuple t is an instance of T .

$$T = (T_1, T_2, \dots, T_n), n \in \mathbb{N} \quad (26.28)$$

$$t = (t_1, t_2, \dots, t_n) \in T \Leftrightarrow t_i \in T_i \forall 0 \leq i < n \wedge \text{len}(t) = \text{len}(T) \quad (26.29)$$

¹¹ http://en.wikipedia.org/wiki/Cartesian_product [accessed 2007-07-03]

¹² http://en.wikipedia.org/wiki/Countable_set [accessed 2007-07-03]

¹³ see definition of function on page 414

¹⁴ http://en.wikipedia.org/wiki/Axiom_of_power_set [accessed 2007-07-03]

¹⁵ <http://en.wikipedia.org/wiki/Tuple> [accessed 2007-07-03]

26.6 Lists

Definition 26.13 (List). Lists¹⁶ are abstract data types which can be regarded as special tuples. They are sequences where every item is of the same type. We introduce functions that will add elements to or remove elements from lists; that sort lists or search within them. Like tuples, lists can be defined using parenthesis in this book. The single elements of a list are accessed by their index written in brackets $((a, b, c) [1] = b)$ where the first element has the index 0 and the last element has the index $n - 1$ (while n is the count of elements in the list: $n = \text{len}((a, b, c)) = 3$). The empty list is abbreviated with $()$.

Definition 26.14 (createList). The $l = \text{createList}(n, q)$ method creates a new list l of the length n filled with the item q .

$$l = \text{createList}(n, q) \Leftrightarrow \text{len}(l) = n \wedge \forall 0 \leq i < n \Rightarrow l[i] = q \quad (26.30)$$

Definition 26.15 (insertListItem). The function $m = \text{insertListItem}(l, i, q)$ creates a new list m by inserting one element q in a list l at the index $0 \leq i \leq \text{len}(l)$. By doing so, it shifts all elements located at index i and above to the right by one position.

$$\begin{aligned} m = \text{insertListItem}(l, i, q) \Leftrightarrow & \text{len}(m) = \text{len}(l) + 1 \wedge m[i] = q \wedge \\ & \forall j : 0 \leq j < i \Rightarrow m[j] = l[j] \wedge \\ & \forall j : i \leq j < \text{len}(l) \Rightarrow m[j+1] = l[j] \end{aligned} \quad (26.31)$$

Definition 26.16 (addListItem). The addListItem function is a shortcut for inserting one item at the end of a list:

$$\text{addListItem}(l, q) \equiv \text{insertListItem}(l, \text{len}(l), q) \quad (26.32)$$

Definition 26.17 (deleteListItem). The function $m = \text{deleteListItem}(l, i)$ creates a new list m by removing the element at index $0 \leq i < \text{len}(l)$ from the list l ($\text{len}(l) \geq i + 1$).

$$\begin{aligned} m = \text{deleteListItem}(l, i) \Leftrightarrow & \text{len}(m) = \text{len}(l) - 1 \wedge \\ & \forall j : 0 \leq j < i \Rightarrow m[j] = l[j] \wedge \\ & \forall j : i < j < \text{len}(l) \Rightarrow m[j-1] = l[j] \end{aligned} \quad (26.33)$$

Definition 26.18 (deleteListRange). The method $m = \text{deleteListRange}(l, i, c)$ creates a new list m by removing c elements beginning at index $0 \leq i < \text{len}(l)$ from the list l ($\text{len}(l) \geq i + c$).

$$\begin{aligned} m = \text{deleteListRange}(l, i, c) \Leftrightarrow & \text{len}(m) = \text{len}(l) - c \wedge \\ & \forall j : 0 \leq j < i \Rightarrow m[j] = l[j] \wedge \\ & \forall j : i + c \leq j < \text{len}(l) \Rightarrow m[j-c] = l[j] \end{aligned} \quad (26.34)$$

Definition 26.19 (appendList). The function $\text{appendList}(l_1, l_2)$ is a shortcut for adding all the elements of a list l_2 to a list l_1 . We define it recursively as:

$$\text{appendList}(l_1, l_2) \equiv \begin{cases} l_1 & \text{if } \text{len}(l_2) = 0 \\ \text{appendList}(\text{addListItem}(l_1, l_2[0]), \text{deleteListItem}(l_2, 0)) & \text{otherwise} \end{cases} \quad (26.35)$$

Definition 26.20 (countItemOccurrences). The function $\text{countItemOccurrences}(x, l)$ returns the number of occurrences of the element x in the list l .

$$\text{countItemOccurrences}(x, l) = |\{i \in 0 \dots \text{len}(l) - 1 : l[i] = x\}| \quad (26.36)$$

¹⁶ http://en.wikipedia.org/wiki/List_%28computing%29 [accessed 2007-07-03]

Definition 26.21 (subList). The method $\text{subList}(l, i, c)$ extracts c elements from the list l beginning at index i and returns them as a new list.

$$\text{subList}(l, i, s) \equiv \text{deleteListRange}(\text{deleteListRange}(l, 0, i), c, |l| - i - c) \quad (26.37)$$

Definition 26.22 (Sorting Lists). It is often useful to have sorted lists¹⁷. Thus we define the functions $S = \text{sortList}_a(U, \text{cmp})$ and $S = \text{sortList}_d(U, \text{cmp})$ which sort a list U in ascending or descending order using a comparator function $\text{cmp}(u_1, u_2)$.

$$S = \text{sortList}_a(U, \text{cmp}) \quad (26.38)$$

$$\forall u \in U \exists i \in [0, \text{len}(U) - 1] : S[i] = u \quad (26.39)$$

$$\text{len}(S) = \text{len}(U) \quad (26.40)$$

$$\forall 0 \leq i < \text{len}(U) - 1 \Rightarrow \text{cmp}(S[i], S[i+1], \leq) 0 \quad (26.41)$$

For $S = \text{sortList}_d(U, \text{cmp})$, only Equation 26.41 changes, the rest stays valid:

$$S = \text{sortList}_d(U, s) \quad (26.42)$$

$$\forall 0 \leq i < \text{len}(U) - 1 \Rightarrow \text{cmp}(S[i], S[i+1], \geq) 0 \quad (26.43)$$

The concept of comparator functions has been introduced in Definition 1.15 on page 19. $\text{cmp}(u_1, u_2)$ returns a negative value if u_1 is smaller than u_2 , a positive number if u_1 is greater than u_2 , and 0 if both are equal. Comparator functions are very versatile, they form the foundation of the sorting mechanisms of the Java framework [1693, 1692], for instance. In global optimization, they are perfectly suited to represent the Pareto dominance or prevalence relations introduced in Section 1.2.4 on page 13 and Section 1.2.7. Sorting according to a specific function f of only one parameter can easily be performed by building the comparator $\text{cmp}(u_1, u_2) \equiv (f(u_1) - f(u_2))$. Thus, we will furthermore synonymously use the sorting predicate also with unary functions f .

$$\text{sortList}(U, f) \equiv \text{sortList}(U, \text{cmp}(u_1, u_2) \equiv (f(u_1) - f(u_2))) \quad (26.44)$$

A list U can be sorted in $\mathbf{O}(\text{len}(U) \log \text{len}(U))$ time complexity. For concrete examples of sorting algorithms, see [1705, 1449, 1706].

Definition 26.23 (Searching in Unsorted Lists). Searching an element u in an unsorted list U means walking through it until either the element is found or the end of the whole list has been scanned, which corresponds to complexity $\mathbf{O}(\text{len}(U))$.

$$\text{searchItem}_u(u, U) = \begin{cases} i : U[i] = u & \text{if } u \in U \\ -1 & \text{otherwise} \end{cases} \quad (26.45)$$

Definition 26.24 (Searching in Sorted Lists). Searching an element s in sorted list S means to perform a binary search¹⁸ returning the index of the element if it is contained in S . If $s \notin S$, a negative number is returned indicating the position where the element could be inserted into the list without violating its order. The function $\text{searchItem}_{as}(s, S)$ searches in an ascending sorted list, $\text{searchItem}_{ds}(s, S)$ searches in a descending sorted list. Searching in a sorted list is done in $\mathbf{O}(\log \text{len}(S))$ time. For concrete algorithm examples, again see [1705, 1449, 1706].

$$\text{searchItem}_{as}(s, S) = \begin{cases} i : S[i] = s & \text{if } s \in S \\ (-i - 1) : (\forall j \geq 0, j < i \Rightarrow S[j] \leq s) \wedge (\forall j < \text{len}(S), j \geq i \Rightarrow S[j] > s) & \text{otherwise} \end{cases} \quad (26.46)$$

$$\text{searchItem}_{ds}(s, S) = \begin{cases} i : S[i] = s & \text{if } s \in S \\ (-i - 1) : (\forall j \geq 0, j < i \Rightarrow S[j] \geq s) \wedge (\forall j < \text{len}(S), j \geq i \Rightarrow S[j] < s) & \text{otherwise} \end{cases} \quad (26.47)$$

¹⁷ http://en.wikipedia.org/wiki/Sorting_algorithm [accessed 2007-07-03]

¹⁸ http://en.wikipedia.org/wiki/Binary_search [accessed 2007-07-03]

Definition 26.25 (*removeListItem*). The function `removeListItem( $l, q$ )` finds one occurrence of an element q in a list l by using the appropriate search algorithm and deletes it (returning a new list m).

$$m = \text{removeListItem}(l, q) \Leftrightarrow \begin{cases} l & \text{if } \text{searchItem}(q, l) < 0 \\ \text{deleteListItem}(l, \text{searchItem}(q, l)) & \text{otherwise} \end{cases} \quad (26.48)$$

We can further define transformations between sets and lists which will implicitly be used when needed in this book. It should be noted that `setToList` is not the inverse function of `listToSet`.

$$\begin{aligned} B = \text{setToList}(\text{set } A) &\Rightarrow \forall a \in A \exists i : B[i] = a \wedge \\ &\forall i \in [0, \text{len}(B) - 1] \Rightarrow B[i] \in A \wedge \\ &\text{len}(\text{setToList}(A)) = |A| \end{aligned} \quad (26.49)$$

$$\begin{aligned} A = \text{listToSet}(\text{tuple } B) &\Rightarrow \forall i \in [0, \text{len}(B) - 1] \Rightarrow B[i] \in A \wedge \\ &\forall a \in A \exists i : \text{len}(B) i = a \wedge \\ &|\text{listToSet}(B)| \leq \text{len}(B) \end{aligned} \quad (26.50)$$

26.7 Binary Relations

Definition 26.26 (Binary Relation). A binary¹⁹ relation²⁰ R is defined as an ordered triple (A, B, P) where A and B are arbitrary sets, and P is a subset of the Cartesian product $A \times B$ (see Equation 26.25). The sets A and B are called the domain and codomain of the relation and P is called its graph. The statement $(a, b) \in P : a \in A \wedge b \in B$ is read “ a is R -related to b ” and is written as $R(a, b)$. The order of the elements in each pair of P is important: If $a \neq b$, then $R(a, b)$ and $R(b, a)$ both can be **true** or **false** independently of each other.

Some types and possible properties of binary relations are listed below and illustrated in Figure 26.2. A binary relation can be [1707]:

- Left-total if

$$\forall a \in A \exists b \in B : R(a, b) \quad (26.51)$$

- Surjective²¹ or right-total if

$$\forall b \in B \exists a \in A : R(a, b) \quad (26.52)$$

- Injective²² if

$$\forall a_1, a_2 \in A, b \in B : R(a_1, b) \wedge R(a_2, b) \Rightarrow a_1 = a_2 \quad (26.53)$$

- Functional if

$$\forall a \in A, b_1, b_2 \in B : R(a, b_1) \wedge R(a, b_2) \Rightarrow b_1 = b_2 \quad (26.54)$$

- Bijective²³ if it is left-total, right-total and functional.
- Transitive²⁴ if

$$\forall a, c \in A, b, c \in B : R(a, c) \wedge R(c, b) \Rightarrow R(a, b) \quad (26.55)$$

¹⁹ http://en.wikipedia.org/wiki/Binary_relation [accessed 2007-07-03]


²⁰ http://en.wikipedia.org/wiki/Relation_%28mathematics%29 [accessed 2007-07-03]

²¹ <http://en.wikipedia.org/wiki/Surjective> [accessed 2007-07-03]

²² <http://en.wikipedia.org/wiki/Injective> [accessed 2007-07-03]

²³ <http://en.wikipedia.org/wiki/Bijective> [accessed 2007-07-03]

²⁴ http://en.wikipedia.org/wiki/Transitive_relation [accessed 2007-07-03]

Fig. 26.2: Properties of a binary relation R with domain A and codomain B .

26.7.1 Functions

Definition 26.27 (Function). A function f is a binary relation with the property that for an element x of the domain²⁵ X there is no more than one element y in the codomain Y such that x is related to y . This uniquely determined element y is denoted by $f(x)$. In other words, a function is a functional binary relation and we can write:

$$\forall x \in X, y_1, y_2 \in Y : f(x, y_1) \wedge f(x, y_2) \Rightarrow y_1 = y_2 \quad (26.56)$$

A function maps each element of X to one element in Y . The domain X is the set of possible input values of f and the codomain Y is the set its *possible* outputs. The set of all *actual* outputs $\{f(x) : x \in X\}$ is called range. This distinction between range and codomain can be made obvious with a small example. The sine function can be defined as a mapping from the real numbers to the real numbers $\sin : \mathbb{R} \mapsto \mathbb{R}$, making $\mathbb{R}$ its codomain. Its actual range however is just the real interval $[-1, 1]$.

Monotonicity

Real functions are monotone, i. e., have the property of monotonicity²⁶, if they preserve a given order²⁷.

Definition 26.28 (Monotonically Increasing). A function $f : X \mapsto Y$ that maps a subset of the real numbers $X \subseteq \mathbb{R}$ to a subset of the real numbers $Y \subseteq \mathbb{R}$ is called monotonic, monotonically increasing, increasing, or non-decreasing, if and only if Equation 26.57 holds.

$$\forall x_1 < x_2, x_1, x_2 \in X \Rightarrow f(x_1) \leq f(x_2) \quad (26.57)$$

²⁵ http://en.wikipedia.org/wiki/Domain_%28mathematics%29 [accessed 2007-07-03]

²⁶ http://en.wikipedia.org/wiki/Monotonic_function [accessed 2007-08-08]

²⁷ Order relations are discussed in Section 26.7.2 on the next page.

Definition 26.29 (Monotonically Decreasing). A function $f : X \mapsto Y$ that maps a subset of the real numbers $X \subseteq \mathbb{R}$ to a subset of the real numbers $Y \subseteq \mathbb{R}$ is called monotonically decreasing, decreasing, or non-increasing, if and only if Equation 26.58 holds.

$$\forall x_1 < x_2, x_1, x_2 \in X \Rightarrow f(x_1) \geq f(x_2) \quad (26.58)$$

26.7.2 Order Relations

All of us have learned the meaning and the importance of order since the earliest years in school. The alphabet is ordered, the natural numbers are ordered, the marks on our school reports are ordered, and so on. Matter of fact, we come into contact with orders even way before entering school by learning to distinguish things according to their size, for instance.

Order relations²⁸ are another type of binary relations which is used to express the order amongst the elements of a set A . Since order relations are imposed on single sets, both their domain and their codomain are the same (A , in this case). For such relations, we can define an additional number of properties which can be used to characterize and distinguish the different types of order relations:

- Antisymmetric:

$$R(a_1, a_2) \wedge R(a_2, a_1) \Rightarrow a_1 = a_2 \quad \forall a_1, a_2 \in A \quad (26.59)$$

- Asymmetric

$$R(a_1, a_2) \Rightarrow \neg R(a_2, a_1) \quad \forall a_1, a_2 \in A \quad (26.60)$$

- Reflexivenss

$$R(a, a) \quad \forall a \in A \quad (26.61)$$

- Irreflexivenss

$$\nexists a \in A : R(a, a) \quad (26.62)$$

All order relations are transitive²⁹, and either antisymmetric or symmetric and either reflexive or irreflexive:

Definition 26.30 (Partial Order). A binary relation R defines a (*non-strict, reflexive*) partial order if and only if it is reflexive, antisymmetric, and transitive.

The $\leq$ and $\geq$ operators, for instance, represent non-strict partial orders on the set of the complex numbers $\mathbb{C}$. Partial orders that correspond to the $>$ and $<$ comparators are called *strict*. The Pareto dominance relation introduced in Definition 1.13 on page 13 is another example for such a strict partial order.

Definition 26.31 (Strict Partial Order). A binary relation R defines a *strict* (or *irreflexive*) partial order if it is irreflexive, asymmetric, and transitive.

Definition 26.32 (Total Order). A total order³⁰ (or linear order, simple order) R on the set A is a partial order which is complete/total.

$$R(a_1, a_2) \vee R(a_2, a_1) \quad \forall a_1, a_2 \in A \quad (26.63)$$

The real numbers $\mathbb{R}$ for example are totally ordered whereas on the set of complex numbers $\mathbb{C}$, only (strict or reflexive) partial (non-total) orders can be defined because it is continuous in two dimensions.

²⁸ http://en.wikipedia.org/wiki/Order_relation [accessed 2007-07-03]

²⁹ See Equation 26.55 on page 413 for the definition of transitivity.

³⁰ http://en.wikipedia.org/wiki/Total_order [accessed 2007-07-03]

26.7.3 Equivalence Relations

Another important class of relations are equivalence relations³¹ [1708, 1709] which are often abbreviated with $\equiv$ or $\sim$, i. e., $a_1 \equiv a_2$ and $a_1 \sim a_2$ mean $R(a_1, a_2)$ for the equivalence relation R imposed on the set A and $a_1, a_2 \in A$. Unlike order relations, equivalence relations are *symmetric*, i. e.,

$$R(a_1, a_2) \Rightarrow R(a_2, a_1) \quad \forall a_1, a_2 \in A \quad (26.64)$$

Definition 26.33 (Equivalence Relation). The binary relation R defines an equivalence relation on the set A if and only if it is reflexive, symmetric, and transitive.

Definition 26.34 (Equivalence Class). If an equivalence relation R is defined on a set A , the subset $A' \subseteq A$ of A is an equivalence class³² if and only if $\forall a_1, a_2 \in A' \Rightarrow R(a_1, a_2)$ ($a_1 \sim a_2$).

³¹ http://en.wikipedia.org/wiki/Equivalence_relation [accessed 2007-07-28]

³² http://en.wikipedia.org/wiki/Equivalence_class [accessed 2007-07-28]

Stochastic Theory and Statistics

In this chapter we give a rough introduction into stochastic theory¹ [1710, 1711, 1712, 1713], which subsumes

- probability² theory³, the mathematical study of phenomena characterized by randomness or uncertainty, and
- statistics⁴, the art of collecting, analyzing, interpreting, and presenting data.

27.1 Probability

Probability theory is used to determine the likeliness of the occurrence of an event under ideal mathematical conditions. [1714, 1715]

Definition 27.1 (Random Experiment). Random experiments can be repeated arbitrary often, their results cannot be predicted.

Definition 27.2 (Elementary Event). The possible outcomes of random situations are called elementary events or samples ω .

Definition 27.3 (Sample Space). The set of all possible outcomes (elementary events, samples) of a random situation is the sample space $\Omega = \{\omega_i : i \in 1..N = |\Omega|\}$.

When throwing dice⁵, for example, the sample space will be $\Omega = \omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6$ whereas ω_i means that the number i was thrown.

Definition 27.4 (Random Event). A random event A is a subset of the sample space Ω ($A \subseteq \Omega$). If $\omega \in A$ occurs, then A is occurs too.

Definition 27.5 (Certain Event). The certain event is the random event will always occur in each repetition of a random experiment. Therefore, it is equal to the whole sample space Ω .

Definition 27.6 (Impossible Event). The impossible event will never occur in any repetition of a random experiment, it is defined as $\emptyset$.

Definition 27.7 (Conflicting Events). Two conflicting events A_1 and A_2 can never occur together in a random experiment. Therefore, $A_1 \cap A_2 = \emptyset$.

¹ <http://en.wikipedia.org/wiki/Stochastic> [accessed 2007-07-03]

² <http://en.wikipedia.org/wiki/Probability> [accessed 2007-07-03]

³ http://en.wikipedia.org/wiki/Probability_theory [accessed 2007-07-03]

⁴ <http://en.wikipedia.org/wiki/Statistics> [accessed 2007-07-03]

⁵ Throwing a dice is discussed as example for stochastic extensively in Section 27.5 on page 448.

27.1.1 Probably as defined by Bernoulli (1713)

In some idealized situations, like throwing ideal coins or ideal dice, all elementary events of the sample space have the same probability de Laplace [1716].

$$P(\omega) = \frac{1}{|\Omega|} \quad \forall \omega \in \Omega \quad (27.1)$$

Equation 27.1 is also called the Laplace-assumption. If it holds, the probability of an event A can be defined as:

$$P(A) = \frac{\text{number of possible events in favor of } A}{\text{number of possible events}} = \frac{|A|}{|\Omega|} \quad (27.2)$$

For many random experiments of this type, we can use combinatorial⁶ approaches in order to determine the number of possible outcomes. Therefore, we want to shortly outline the mathematical concepts of factorial numbers, combinations, and permutations.⁷

Definition 27.8 (Factorial). The factorial⁸ $n!$ of a number $n \in \mathbb{N}_0$ is the product of n and all natural numbers smaller than it. It is a specialization of the Gamma function for positive integer numbers, see Section 27.8.1 on page 466.

$$n! = \prod_{i=1}^n i \quad (27.3)$$

$$0! = 1 \quad (27.4)$$

In combinatorial mathematics⁹, we often want to know in how many ways we can arrange $n \in \mathbb{N}$ elements from a set Ω with $M = |\Omega| \geq n$ elements. We can distinguish between *combinations*, where the order of the elements in the arrangement plays no role, and *permutations*, where it is important. (a, b, c) and (c, b, a) , for instance, denote the same combination but different permutations of the elements $\{a, b, c\}$. We furthermore distinguish between arrangements where each element of Ω can occurred at most once (without repetition) and arrangements where the same elements may occur multiple time (with repetition).

Combinations

The number of possible combinations¹⁰ $C(M, n)$ of $n \in \mathbb{N}$ elements out of a set Ω with $M = |\Omega| \geq n$ elements without repetition is

$$C(M, n) = \binom{M}{n} = \frac{M!}{n!(M-n)!} \quad (27.5)$$

$$\binom{M}{n} = \frac{M}{n} * \frac{M-1}{n-1} * \frac{M-2}{n-2} * \dots * \frac{M-n+1}{1} \quad (27.6)$$

$$C(M+1, n) = C(M, n) + C(M, n-1) = \binom{M+1}{n} = \binom{M}{n} + \binom{M}{n-1} \quad (27.7)$$

If the elements of Ω may repeatedly occur in the arrangements, the number of possible combinations becomes

$$\frac{M+n-1!}{n!(M-1)!} = \binom{M+n-1}{n} = \binom{M+n-1}{n-1} = C(M+n-1, n) = C(M+n-1, n-1) \quad (27.8)$$

⁶ <http://en.wikipedia.org/wiki/Combinatorics> [accessed 2007-07-03]

⁷ http://en.wikipedia.org/wiki/Combinations_and_permutations [accessed 2007-07-03]

⁸ <http://en.wikipedia.org/wiki/Factorial> [accessed 2007-07-03]

⁹ <http://en.wikipedia.org/wiki/Combinatorics> [accessed 2008-01-31]

¹⁰ <http://en.wikipedia.org/wiki/Combination> [accessed 2008-01-31]

Permutations

The number of possible permutations¹¹ $Perm(M, n)$ of $n \in \mathbb{N}$ elements out of a set Ω with $M = |\Omega| \geq n$ elements without repetition is

$$Perm(M, n) = (M)_n = \frac{M!}{(M - n)!} \quad (27.9)$$

If an element from Ω can occur more than once in the arrangements, the number of possible permutations is

$$M^n \quad (27.10)$$

27.1.2 The Limiting Frequency Theory of von Mises

If we repeat a random experiment multiple times, the number of occurrences of a certain event should somehow reflect its probability. The more often we perform the experiment, the more reliable will the estimations of the event probability become. We can express this relation using the notation of *frequency*.

Definition 27.9 (Absolute Frequency). The number $H(A, n)$ denoting how often an event A occurred during n repetitions of a random experiment is its absolute frequency¹².

Definition 27.10 (Relative Frequency). The relative frequency $h(A, n)$ of an event A is its absolute frequency normalized to the total number of experiments n . The relative frequency has the following properties:

$$h(A, n) = \frac{H(A, n)}{n} \quad (27.11)$$

$$0 \leq h(A, n) \leq 1 \quad (27.12)$$

$$h(\Omega, n) = 1 \quad \forall n \in \mathbb{N} \quad (27.13)$$

$$A \cap B = \emptyset \Rightarrow h(A \cup B, n) = \frac{H(A, n) + H(B, n)}{n} = h(A, n) + h(B, n) \quad (27.14)$$

According to von Mises [1717], the (statistical) probability $P(A)$ of an event A computing the limit of its relative frequency $h(A, n)$ as n approaching infinity. This is the limit of the quotient of the number of elementary events favoring A and the number of all possible elementary events for infinite many repetitions. [1717, 1718]

$$P(A) = \lim_{n \rightarrow \infty} h(A, n) = \lim_{n \rightarrow \infty} \frac{n_A}{n} \quad (27.15)$$

27.1.3 The Axioms of Kolmogorov

Definition 27.11 (σ -algebra). A subset S of the power set $\mathcal{P}(\Omega)$ is called σ -algebra¹³, if the following axioms hold:

$$\Omega \in S \quad (27.16)$$

$$\emptyset \in S \quad (27.17)$$

$$A \in S \Leftrightarrow \bar{A} \in S \quad (27.18)$$

$$A \in S \wedge B \in S \Rightarrow (A \cup B) \in S \quad (27.19)$$

¹¹ <http://en.wikipedia.org/wiki/Permutations> [accessed 2008-01-31]

¹² http://en.wikipedia.org/wiki/Frequency_%28statistics%29 [accessed 2007-07-03]

¹³ <http://en.wikipedia.org/wiki/Sigma-algebra> [accessed 2007-07-03]

From these axioms others can be deduced, for example:

$$A \in S \wedge B \in S \Rightarrow \bar{A} \in S \wedge \bar{B} \in S \quad (27.20)$$

$$\begin{aligned} &\Rightarrow \overline{A \cup B} \in S \\ &\Rightarrow \overline{A \cap B} \in S \end{aligned} \quad (27.21)$$

$$A \in S \wedge B \in S \Rightarrow (A \cap B) \in S \quad (27.22)$$

Definition 27.12 (Probability Space). A probability space (or random experiment) is defined by the triple (Ω, S, P) whereas

- Ω is the sample space, a set of elementary events,
- S is a σ -algebra defined on Ω , and
- P defines a probability measure¹⁴ that determines the probability of occurrence for each event $\omega \in \Omega$. (Kolmogorov [1719] axioms¹⁵)

Definition 27.13 (Probability). A mapping P which maps a real number to each elementary event $\omega \in \Omega$ is called probability measure if and only if the σ -algebra S on Ω holds:

$$\forall A \in S \Rightarrow 0 \leq P(A) \leq 1 \quad (27.23)$$

$$P(\Omega) = 1 \quad (27.24)$$

$$\forall \text{disjoint } A_i \in S \Rightarrow P(A) = P\left(\bigcup_{\forall i} A_i\right) = \sum_{\forall i} P(A_i) \quad (27.25)$$

From this axioms can be deduced:

$$P(\emptyset) = 0 \quad (27.26)$$

$$P(A) = 1 - P(\bar{A}) \quad (27.27)$$

$$P(A \cap \bar{B}) = P(A) - P(A \cap B) \quad (27.28)$$

$$P(A \cup B) = P(A) + P(B) - P(A \cap B) \quad (27.29)$$

27.1.4 Conditional Probability

Definition 27.14 (Conditional Probability). The conditional probability¹⁶ $P(A|B)$ is the probability of some event A , given the occurrence of some other event B . $P(A|B)$ is read “the probability of A , given B ”.

$$P(A|B) = \frac{P(A \cap B)}{P(B)} \quad (27.30)$$

$$P(A \cap B) = P(A|B) P(B) \quad (27.31)$$

Definition 27.15 (Statistical Independence). Two events A and B are (statistically) independent if and only if $P(A \cap B) = P(A) P(B)$ holds. If two events A and B are statistically independent, we can deduce:

$$P(A \cap B) = P(A) P(B) \quad (27.32)$$

$$P(A|B) = P(A) \quad (27.33)$$

$$P(B|A) = P(B) \quad (27.34)$$

¹⁴ http://en.wikipedia.org/wiki/Probability_measure [accessed 2007-07-03]

¹⁵ http://en.wikipedia.org/wiki/Kolmogorov_axioms [accessed 2007-07-03]

¹⁶ http://en.wikipedia.org/wiki/Conditional_probability [accessed 2007-07-03]

27.1.5 Random Variable

Definition 27.16 (Random Variable). The function X which relates the sample space Ω to the real numbers $\mathbb{R}$ is called random variable¹⁷ in the probability space (Ω, S, P) .

$$X : \Omega \mapsto \mathbb{R} \quad (27.35)$$

Using such a random variable, we can replace the sample space Ω with the new sample space Ω_X . Furthermore, the σ -algebra S can be replaced with a σ -algebra S_X , which consists of subsets of Ω_X instead of Ω . Last but not least, we replace the probability measure P which relates the $\omega \in \Omega$ to the interval $[0, 1]$ by a new probability measure P_X which relates the real numbers $\mathbb{R}$ to this interval.

Definition 27.17 (Probability Space of a Random Variable). Is $X : \Omega \mapsto \mathbb{R}$ a random variable, then the probability space of X is defined as the triplet

$$(\Omega_X, S_X, P_X) \quad (27.36)$$

One example for such a new probability measure would be the probability that a random variable X takes on a real value which is smaller or equal a value x :

$$P_X(X \leq x) = P(\{\omega : \omega \in \Omega \wedge X(\omega) \leq x\}) \quad (27.37)$$

27.1.6 Cumulative Distribution Function

Definition 27.18 (Cumulative Distribution Function). If X is a random variable of a probability space $(\Omega_X = \mathbb{R}, S_X, P_X)$, we call the function $F_X : \mathbb{R} \mapsto [0, 1]$ the (cumulative) distribution function¹⁸ (CDF) of the random variable X if it fulfills Equation 27.38.

$$F_X := \underbrace{P_X(X \leq x)}_{\text{definition rnd. var.}} \equiv \underbrace{P(\{\omega : \omega \in \Omega \wedge X(\omega) \leq x\})}_{\text{definition probability space}} \quad (27.38)$$

A cumulative distribution function F_X has the following properties:

- $F_X(X)$ is normalized:

$$\underbrace{\lim_{x \rightarrow -\infty} F_X(x) = 0}_{\text{impossible event}}, \quad \underbrace{\lim_{x \rightarrow +\infty} F_X(x) = 1}_{\text{certain event}} \quad (27.39)$$

- $F_X(X)$ is monotonously¹⁹ growing:

$$F_X(x_1) \leq F_X(x_2) \quad \forall x_1 \leq x_2 \quad (27.40)$$

- $F_X(X)$ is (right-sided) continuous²⁰:

$$\lim_{h \rightarrow 0} F_X(x + h) = F_X(x) \quad (27.41)$$

- The probability that the random variable X takes on values in the interval $x_0 \leq X \leq x_1$ can be computed using the CDF:

$$P(x_0 \leq X \leq x_1) = F_X(x_1) - F_X(x_0) \quad (27.42)$$

¹⁷ http://en.wikipedia.org/wiki/Random_variable [accessed 2007-07-03]

¹⁸ http://en.wikipedia.org/wiki/Cumulative_distribution_function [accessed 2007-07-03]

¹⁹ <http://en.wikipedia.org/wiki/Monotonicity> [accessed 2007-07-03]

²⁰ http://en.wikipedia.org/wiki/Continuous_function [accessed 2007-07-03]

- The probability that the random variable X takes on the value of a single random number x :

$$P(X = x) = F_X(x) - \lim_{h \rightarrow 0} F_X(x - h) \quad (27.43)$$

We can further distinguish between sample spaces Ω which contain at most countable infinite many elements and such that are continuums. Hence, we there are discrete²¹ and continuous²² random variables.

Definition 27.19 (Discrete Random Variable). A random variable X (and its probability measure $P_X(X)$ respectively) is called discrete if it takes on at most countable infinite many values. Its cumulative distribution function $F_X(X)$ therefore has the shape of a stair-way.

Definition 27.20. A random variable X (and its probability measure P_X respectively) is called continuous if it can take on uncountable infinite many values and its cumulative distribution function $F_X(X)$ is also continuous.

27.1.7 Probability Mass Function

Definition 27.21 (Probability Mass Function). The probability mass function²³ (PMF) f_X is defined discrete distributions only and assigns a probability to each value a discrete random variable X can take on.

$$f_X : \mathbb{Z} \mapsto [0, 1] : f_X(x) := P_X(X = x) \quad (27.44)$$

Therefore, we can specify the relation between the PMF and its corresponding (discrete) CDF in Equation 27.45 and Equation 27.45. We can further define the probability of an event A in Equation 27.47 using the PMF.

$$P_X(X \leq x) = F_X(x) = \sum_{i=-\infty}^x f_X(x) \quad (27.45)$$

$$P_X(X = x) = f_X(x) = F_X(x) - F_X(x - 1) \quad (27.46)$$

$$P_X(A) = \sum_{\forall x \in A} f_X(x) \quad (27.47)$$

27.1.8 Probability Density Function

The probability density function²⁴ (PDF) is the counterpart of the PMF for continuous distributions. The PDF *does not* represent the probabilities of the single values of a random variable. Since a continuous random variable can take on uncountable many values, each distinct value itself has the probability 0. If we, for instance, picture the current temperature outside as (continuous) random variable, the probability that it takes on the value 18 for 18°C is zero. It will never be *exactly* 18°C outside, we can at most declare with a certain probability that we have a temperature between 17.99999°C and 18.00001°C.

Definition 27.22 (Probability Density Function). If a random variable X is continuous, its probability density function f_X is defined as

$$f_X : \mathbb{R} \mapsto [0, \infty) : F_X(x) = \int_{-\infty}^{+\infty} f_X(\xi) d\xi \quad \forall x \in \mathbb{R} \quad (27.48)$$

²¹ http://en.wikipedia.org/wiki/Discrete_random_variable [accessed 2007-07-03]

²² http://en.wikipedia.org/wiki/Continuous_probability_distribution [accessed 2007-07-03]

²³ http://en.wikipedia.org/wiki/Probability_mass_function [accessed 2007-07-03]

²⁴ http://en.wikipedia.org/wiki/Probability_density_function [accessed 2007-07-03]

27.2 Stochastic Properties

Each random variable X which conforms to a probability distribution F_X may have certain properties such as a maximum and a mean value, a variance, and a value which will be taken on by X most often. If the cumulative distribution function F_X of X is known, these values can usually be computed directly from its parameters. On the other hand, it is possible that we only know the values $A[i]$ which X took on during some random experiments. From this set of sample data A , we can *estimate* the properties of the underlying (possibly unknown) distribution of X using statistical methods (with a certain error, of course).

In the following, we will elaborate on the properties of a random variable $X \in \mathbb{R}$ both from the viewpoint of knowing the PMF/PDF $f_X(x)$ and the CDF $F_X(x)$ as well as from the statistical perspective, where only a sample A of past values of X is known. In the latter case, we define the sample as a list A with the length $n = \text{len}(A)$ and the elements $A[i] : i \in [0, n - 1]$.

27.2.1 Count, Min, Max and Range

The most primitive features of a random distribution are the minimum, maximum, and the range of its values, as well as the number of values $A[i]$ in a data sample A .

Definition 27.23 (Count). $n = \text{len}(A)$ is the number of elements in the data sample A .

This item count is only defined for data samples, not for random variables, since random variables represent experiments which can infinitely be repeated and thus stand for infinitely many values. The number of items should not be mixed up with the possible number of different values the random variable may take on. A data sample A may contain the same value a multiple times. When throwing a dice seven times, one may throw $A = (1, 4, 3, 3, 2, 6, 1)$, for example²⁵.

Definition 27.24 (Minimum). There exists no smaller element in the sample data A than the minimum sample $\tilde{a} \equiv \min(A)$ when speaking statistics. From the perspective of the cumulative distribution function F_X , the minimum is the lower boundary $\tilde{x}$ of the random variable X (or negative infinity, if no such boundary exists). Both definitions are fully compliant to Definition 1.10 on page 8.

$$\min(A) \equiv \tilde{a} \in A : \forall a \in A \Rightarrow \tilde{a} \leq a \quad (27.49)$$

$$\tilde{x} = \min(X) \Leftrightarrow F_X(\tilde{x}) > 0 \wedge F_X(x) \geq F_X(\tilde{x}) \quad \forall x \in \mathbb{R} \quad (27.50)$$

Definition 27.25 (Maximum). In statistically evaluated sample data, exists element bigger than the maximum $\hat{a} = \max(A)$ in A . The value $\hat{x}$ is the upper boundary of the values a random variable X may take on (or positive infinity, if X is unbounded). This definition is compliant with Definition 1.9 on page 8.

$$\max(A) \equiv \hat{a} \in A : \forall a \in A \Rightarrow \hat{a} \geq a \quad (27.51)$$

$$\hat{x} = \max(X) \Leftrightarrow F_X(\hat{x}) \geq F_X(x) \quad \forall x \in \mathbb{R} \quad (27.52)$$

Definition 27.26 (Range).

The range $\text{range}(A)$ of the sample data A is the difference of the maximum $\max(A)$ and the minimum $\min(A)$ of A and therefore represents the span covered by the data. If a random variable X is limited in both directions, it has a finite range $\text{range}(X)$, otherwise this range is infinite too.

$$\text{range}(A) = \hat{a} - \tilde{a} = \max(A) - \min(A) \quad (27.53)$$

$$\text{range}(X) = \hat{x} - \tilde{x} = \max(X) - \min(X) \quad (27.54)$$

²⁵ Throwing a dice is discussed as example for stochastic extensively in Section 27.5 on page 448.

27.2.2 Expected Value and Arithmetic Mean

The expected value EX and the $\bar{a}$ are basic measures for random variables and data samples that help us to estimate the regions where their values will be distributed around.

Definition 27.27 (Expected Value). The expected value²⁶ of a random variable X is the sum of the probability of each possible outcome of the random experiment multiplied by the outcome value. It is abbreviated by EX or μ . For discrete distributions it can be computed using Equation 27.55 and for continuous ones Equation 27.56 holds.

$$EX = \sum_{x=-\infty}^{\infty} x f_X(x) \quad (27.55)$$

$$EX = \int_{-\infty}^{\infty} x f_X(x) dx \quad (27.56)$$

If the expected value EX of a random variable X is known, the following statements can be derived the expected values of some related random variables as follows:

$$Y = a + X \Rightarrow EY = a + EX \quad (27.57)$$

$$Z = bX \Rightarrow EZ = bEX \quad (27.58)$$

Definition 27.28 (Sum). The $\text{sum}(A)$ represents the sum of all elements in a set of data samples A . This value does, of course, only exist in statistics.

$$\text{sum}(A) = \sum_{i=0}^{n-1} A[i] \quad (27.59)$$

Definition 27.29 (Arithmetic Mean). The arithmetic mean²⁷ $\bar{a}$ is the sum of all elements in the sample data A divided by the total number of values. In the spirit of the limiting frequency method of von Mises [1717], it is an estimation of the expected value $\bar{a} \approx EX$ of the random variable X that produced the sample data A .

$$\bar{a} = \frac{\text{sum}(A)}{n} = \frac{1}{n} \sum_{i=0}^{n-1} A[i] = \sum_{i=0}^{n-1} h(A[i], n) \quad (27.60)$$

27.2.3 Variance and Standard Deviation

The variance²⁸ [182] is a measure of statistical dispersion. It illustrates how close the results of a random variable or the elements a in a data sample A are to their expected value EX or their arithmetical mean $\bar{a}$.

Definition 27.30 (Variance of a Random Variable). The variance $D^2X \equiv \text{var}(X) \equiv \sigma^2$ of a random variable X is defined as

$$\text{var}(X) = D^2X = E[(X - EX)^2] = E[X^2] - (EX)^2 \quad (27.61)$$

The variance of a discrete random variable X can be computed using Equation 27.62 and for continuous distributions, Equation 27.63 will hold.

²⁶ http://en.wikipedia.org/wiki/Expected_value [accessed 2007-07-03]

²⁷ http://en.wikipedia.org/wiki/Arithmetic_mean [accessed 2007-07-03]

²⁸ <http://en.wikipedia.org/wiki/Variance> [accessed 2007-07-03]

$$\begin{aligned}
D^2X &= \sum_{x=-\infty}^{\infty} f_X(x) (x - EX)^2 = \sum_{x=-\infty}^{\infty} x^2 f_X(x) - \left[\sum_{x=-\infty}^{\infty} x f_X(x) \right]^2 \\
&= \left[\sum_{x=-\infty}^{\infty} x^2 f_X(x) \right] - (EX)^2
\end{aligned} \tag{27.62}$$

$$\begin{aligned}
D^2X &= \int_{-\infty}^{\infty} (x - EX)^2 dx = \int_{-\infty}^{\infty} x^2 f_X(x) dx - \left[\int_{-\infty}^{\infty} x f_X(x) dx \right]^2 \\
&= \left[\int_{-\infty}^{\infty} x^2 f_X(x) dx \right] - (EX)^2
\end{aligned} \tag{27.63}$$

If the variance D^2X of a random variable X is known, we can derive the variances of some related random variables as follows:

$$Y = a + X \Rightarrow D^2Y = D^2X \tag{27.64}$$

$$Z = bX \Rightarrow D^2Z = b^2 D^2X \tag{27.65}$$

Definition 27.31 (Sum of Squares). The function $\text{sumSqs}(A)$ is only defined for statistical data and represents the sum of the squares of all elements in the data sample A .

$$\text{sumSqs}(A) = \sum_{i=0}^{n-1} (A[i])^2 \tag{27.66}$$

Definition 27.32 (Variance Estimator). We define the (unbiased) estimator²⁹ s^2 of the variance of the random variable which produced the sample values A according to Equation 27.67. The variance is zero for all samples with $(n = \text{len}(A)) \leq 1$.

$$s^2 = \frac{1}{n-1} \sum_{i=0}^{n-1} (A[i] - \bar{a})^2 = \frac{1}{n-1} \left(\text{sumSqs}(A) - \frac{(\text{sum}(A))^2}{n} \right) \tag{27.67}$$

Definition 27.33 (Standard Deviation). The standard deviation³⁰ is the square root of the variance. The standard deviation of a random variable X is abbreviated with DX and σ , its statistical estimate is s .

$$DX = \sqrt{D^2X} \tag{27.68}$$

$$s = \sqrt{s^2} \tag{27.69}$$

The standard deviation is zero for all samples with $n \leq 1$.

Definition 27.34 (Coefficient of Variation). The coefficient of variation³¹ c_V of a random variable X is the ratio of the standard deviation by expected value of X . For data samples, its estimate $c_{\approx V}$ is defined as the ration of the estimate of the standard deviation and the arithmetic mean.

$$c_V = \frac{DX}{EX} = \frac{\sigma}{\mu} \tag{27.70}$$

$$c_{\approx V} = \frac{n}{\text{sum}(A)} \sqrt{\frac{\text{sumSqs}(A) - \frac{(\text{sum}(A))^2}{n}}{n-1}} \tag{27.71}$$

²⁹ see Definition 27.54 on page 450

³⁰ http://en.wikipedia.org/wiki/Standard_deviation [accessed 2007-07-03]

³¹ http://en.wikipedia.org/wiki/Coefficient_of_variation [accessed 2007-07-03]

Definition 27.35 (Covariance). The covariance³² $\text{cov}(X, Y)$ of two random variables X and Y is a measure for how much they are related. It exists if the expected values EX^2 and EY^2 exist and is defined as

$$\text{cov}(X, Y) = E[X - EX] * E[Y - EY] \quad (27.72)$$

$$= E[X * Y] - EX * EY \quad (27.73)$$

$$(27.74)$$

If X and Y are statistically independent, then their covariance is zero, since

$$E[X * Y] = EX * EY \quad (27.75)$$

Furthermore, the following formulas hold for the covariance

$$D^2 X = \text{cov}(X, X) \quad (27.76)$$

$$D^2[X + Y] = \text{cov}(X + Y, X + Y) = D^2 X + D^2 Y + 2\text{cov}(X, Y) \quad (27.77)$$

$$D^2[X - Y] = \text{cov}(X - Y, X + Y) = D^2 X + D^2 Y - 2\text{cov}(X, Y) \quad (27.78)$$

$$\text{cov}(X, Y) = \text{cov}(Y, X) \quad (27.79)$$

$$\text{cov}(aX, Y) = a \text{cov}(Y, X) \quad (27.80)$$

$$\text{cov}(X + Y, Z) = \text{cov}(X, Z) + \text{cov}(Y, Z) \quad (27.81)$$

$$\text{cov}(aX + b, cY + d) = a c \text{cov}(X, Y) \quad (27.82)$$

27.2.4 Moments

Definition 27.36 (Moment). The k^{th} moment³³ $\mu'_k(c)$ about a value c is defined for a random distribution X as

$$\mu'_k(c) = E[(X - c)^k] \quad (27.83)$$

It can be specified for discrete (Equation 27.84) and continuous (Equation 27.85) probability distributions using Equation 27.55 and Equation 27.56 as follows.

$$\mu'_k(c) = \sum_{x=-\infty}^{\infty} f_X(x) (x - c)^k \quad (27.84)$$

$$\mu'_k(c) = \int_{-\infty}^{\infty} f_X(x) (x - c)^k dx \quad (27.85)$$

Definition 27.37 (Statistical Moment). The k^{th} statistical moment μ'_k of a random distribution is its k^{th} moment about zero, i. e., the expected value of its values raised to the k^{th} power.

$$\mu'_k = \mu'_k(0) = E[X^k] \quad (27.86)$$

Definition 27.38 (Central Moment). The k^{th} moment about the mean (or central moment)³⁴ is the expected value of the difference between elements and their expected value raised to the k^{th} power.

$$\mu_k = E[(X - EX)^k] \quad (27.87)$$

Hence, the variance $D^2 X$ equals the second central moment μ_2 .

Definition 27.39 (Standardized Moment). The k^{th} standardized moment $\mu_{\sigma,k}$ is the quotient of the k^{th} central moment and the standard deviation raised to the k^{th} power.

$$\mu_{\sigma,k} = \frac{\mu_k}{\sigma^k} \quad (27.88)$$

³² <http://en.wikipedia.org/wiki/Covariance> [accessed 2008-02-05]

³³ http://en.wikipedia.org/wiki/Moment_%28mathematics%29 [accessed 2008-02-01]

³⁴ http://en.wikipedia.org/wiki/Moment_about_the_mean [accessed 2007-07-03]

27.2.5 Skewness and Kurtosis

The two other most important moments of random distributions are the skewness γ_1 and the kurtosis γ_2 and their estimates G_1 and G_2 .

Definition 27.40 (Skewness). The skewness³⁵ γ_1 , the third standardized moment, is a measure of asymmetry of a probability distribution. If $\gamma_1 > 0$, the right part of the distribution function is either longer or fatter (positive skew, right-skewed). If $\gamma_1 < 0$, the distribution's left part is longer or fatter.

$$\gamma_1 = \mu_{\sigma,3} = \frac{\mu_3}{\sigma^3} \quad (27.89)$$

For sample data A the skewness of the underlying random variable is approximated with the estimator G_1 where s is the estimated standard deviation. The sample skewness is only defined for sets A with at least three elements.

$$G_1 = \frac{n}{(n-1)(n-2)} \sum_{i=0}^{n-1} \left(\frac{A[i] - \bar{a}}{s} \right)^3 \quad (27.90)$$

Definition 27.41 (Kurtosis). The excess kurtosis³⁶ γ_2 is a measure for the sharpness of a distribution's peak. A distribution with a high kurtosis has a sharper "peak" and fatter "tails", while a distribution with a low kurtosis has a more rounded peak with wider "shoulders". The normal distribution (see Section 27.4.2) has a zero kurtosis.

$$\gamma_2 = \mu_{\sigma,4} - 3 = \frac{\mu_4}{s^3} - 3 \quad (27.91)$$

For sample data A which represents only a subset of a greater amount of data, the sample kurtosis can be approximated with the estimator G_2 where s is the estimate of the sample's standard deviation. The kurtosis is only defined for sets with at least four elements.

$$G_2 = \left[\frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_{i=0}^{n-1} \left(\frac{A[i] - \bar{a}}{s} \right)^4 \right] - \frac{3(n-1)^2}{(n-2)(n-3)} \quad (27.92)$$

27.2.6 Median, Quantiles, and Mode

Definition 27.42 (Median). The median $m = \text{med}(X)$ is the value right in the middle of a sample or distribution, dividing it into two equal halves. Therefore, the probability of drawing an element less than $\text{med}(X)$ is equal to the probability of drawing an element larger than m .

$$P(X \leq m) \geq \frac{1}{2} \wedge P(X \geq m) \geq \frac{1}{2} \wedge P(X \leq m) \leq P(X \geq m) \quad (27.93)$$

We can determine the median m of continuous and discrete distributions by solving Equation 27.94 and Equation 27.95 respectively.

$$\frac{1}{2} = \int_{-\infty}^m f_X(x) dx \quad (27.94)$$

$$\sum_{i=-\infty}^{m-1} f_X(x) \leq \frac{1}{2} \leq \sum_{i=m}^{\infty} f_X(x) \quad (27.95)$$

$$(27.96)$$

³⁵ <http://en.wikipedia.org/wiki/Skewness> [accessed 2007-07-03]

³⁶ <http://en.wikipedia.org/wiki/Kurtosis> [accessed 2008-02-01]

If a sample A has an odd element count, the median m is the element in the middle, otherwise (in a set with an even element count there exists no single “middle”-element), the arithmetic mean of the two middle elements. The median represents the dataset in an unbiased manner. If you have, for example, the dataset $A = (1, 1, 1, 1, 1, 2, 2, 2, 500\,000)$, the arithmetic mean, biased by the large element 500 000 would be very high ($55556.\bar{7}$). The median however would be 1 and thus represents the sample better. The median of a sample can be computed as:

$$A_s \equiv \text{sortList}_a(A, >) \quad (27.97)$$

$$\text{med}(A) = \begin{cases} A_s[\frac{n-1}{2}] & \text{if } n = \text{len}(A) \text{ is odd} \\ \frac{1}{2} (A_s[\frac{n}{2}] + A_s[\frac{n}{2}-1]) & \text{otherwise} \end{cases} \quad (27.98)$$

Definition 27.43 (Quantile). Quantiles³⁷ are points taken at regular intervals from a sorted dataset (or a cumulative distribution function). The q -quantiles divide a distribution function F_X or data sample A into q parts T_i with equal probability. They can be regarded as the generalized median, or vice versa, the median is the 2-quantile.

$$\forall x \in \mathbb{R}, i \in [0, q-1] \Rightarrow \frac{1}{q} \leq P(x \in T_i) \quad (27.99)$$

A sorted data sample is divided into q subsets of equal length by the q -quantiles. The cumulative distribution function of a random variable X is divided by the q -quantiles into q subsets of equal area. The quantiles are the boundaries between the subsets/areas. Therefore, the k^{th} q -quantile is the value ζ so that the probability that the random variable (or an element of the data set) will take on a value less than ζ is at most $\frac{k}{q}$ and the probability that it will take on a value greater than or equal to ζ is at most $\frac{q-k}{q}$. There exist $q-1$ q -quantiles (k spans from 1 to $q-1$). The k^{th} q -quantile $\text{quantile}_q^k(A)$ of a dataset A can be computed as:

$$A_s \equiv \text{sortList}_a(A, >) \quad (27.100)$$

$$\text{quantile}_q^k(A) = A_s[\lfloor \frac{k*n}{q} \rfloor] \quad (27.101)$$

For some special values of q , the quantiles have been given special names too (see Table 27.1).

q name
100 percentiles
10 deciles
9 noniles
5 quintiles
4 quartiles
2 median

Table 27.1: Special Quantiles

Definition 27.44 (Interquartile Range). The interquartile range³⁸ is the range between the first and the third quartile and defined as $\text{quantile}_4^3(X) - \text{quantile}_4^1(X)$.

Definition 27.45 (Mode). The mode³⁹ is the value that most often occurs in a data sample or is most frequently assumed by a random variable. There exist unimodal distributions/samples that have one mode value and multimodal distributions/samples with multiple modes.

³⁷ <http://en.wikipedia.org/wiki/Quantiles> [accessed 2007-07-03]

³⁸ http://en.wikipedia.org/wiki/Inter-quartile_range [accessed 2007-07-03]

³⁹ http://en.wikipedia.org/wiki/Mode_%28statistics%29 [accessed 2007-07-03]

In [1720, 1721] you can find further information of the relation between the mode, the mean and the skewness.

27.2.7 Entropy

Definition 27.46 (Information Entropy). The information entropy⁴⁰ $H(X)$ defined by Shannon [1722] is a measure of uncertainty for discrete probability mass functions f_X of random variables X or data sets A . It is defined in as follows. The $h(a, n)$ in Equation 27.103 denotes the relative frequency of the value a amongst the n samples in A .

$$H(X) = \sum_{x=-\infty}^{\infty} f_X(x) \log_2 \left(\frac{1}{f_X(x)} \right) = - \sum_{x=-\infty}^{\infty} f_X(x) \log_2 f_X(x) \quad (27.102)$$

$$H(A) = - \sum_{\forall a \in A} h(a, n) \log_2 h(a, n) \quad (27.103)$$

Definition 27.47 (Differential Entropy). The differential (also called continuous) entropy $h(X)$ is a generalization of the information entropy to continuous probability density functions f_X of random variables X . [1723]

$$h(X) = - \int_{-\infty}^{\infty} f_X(x) \ln f_X(x) dx \quad (27.104)$$

27.2.8 The Law of Large Numbers

The law of large numbers (LLN) combines statistics and probability by showing that if an event e with the probability $P(e) = p$ is observed in n independent repetitions of a random experiment, its relative frequency $h(e, n)$ (see Definition 27.10) converges to its probability p if n becomes larger.

In the following, assume that the A is an infinite sequence of samples from equally distributed and pairwise independent random variables X_i with the (same) expected value EX . The weak law of large numbers states that the mean $\bar{a}$ of the sequence A converges to a value in $(EX - \varepsilon, EX + \varepsilon)$ for each positive real number $\varepsilon > 0, \varepsilon \in \mathbb{R}^+$.

$$\lim_{n \rightarrow \infty} P(|\bar{a} - EX| < \varepsilon) = 1 \quad (27.105)$$

In other words, the weak law of large numbers says that the sample average will converge to the expected value of the random experiment if the experiment is repeated many times.

According to the strong law of large numbers, the mean $\bar{a}$ of the sequence A even converges to the expected value EX of the underlying distribution for infinite large n

$$P\left(\lim_{n \rightarrow \infty} \bar{a} = EX\right) = 1 \quad (27.106)$$

The law of large numbers implies that the accumulated results of each random experiment will approximate the underlying distribution function if repeated infinitely (under the condition that there exists an invariable underlying distribution function).

27.3 Some Discrete Distributions

In this section we will introduce some common discrete distributions. Discrete probability distributions assign probabilities to the elements of a finite (or, at most, countable infinite) set of discrete events/outcomes of a random experiment.

Parts of the information provided in this and the following section have been obtained from Wikipedia [2].

⁴⁰ http://en.wikipedia.org/wiki/Information_entropy [accessed 2007-07-03]

27.3.1 Discrete Uniform Distribution

The uniform distribution exists in a discrete⁴¹ as well as in a continuous form. In this section we want to discuss the discrete form whereas the continuous form is elaborated on in Section 27.3.1.

All possible outcomes $\omega \in \Omega$ of a random experiment which obeys the uniform distribution have exactly the same probability. In the discrete uniform distribution, Ω has at most countable infinite elements (although normally being finite). The best example for this distribution is throwing an ideal dice. This experiment has six possible outcomes ω_i where each has the same probability $P(\omega_i) = \frac{1}{6}$. Throwing ideal coins and drawing one element out of a set of n possible elements are other examples where a discrete uniform distribution can be assumed. Table 27.2 contains the characteristics of the discrete uniform distribution. In Figure 27.1 you can find some example uniform probability mass functions and in Figure 27.2 we have outlined their according cumulative distribution functions.

parameter	definition	
parameters	$a, b \in \mathbb{Z}, a > b$	(27.107)
$ \Omega $	$ \Omega = r = \text{range} = b - a + 1$	(27.108)
PMF	$P(X = x) = f_X(x) = \begin{cases} \frac{1}{r} & \text{if } a \leq x \leq b, x \in \mathbb{Z} \\ 0 & \text{otherwise} \end{cases}$	(27.109)
CDF	$P(X \leq x) = F_X(x) = \begin{cases} 0 & \text{if } x < a \\ \lfloor \frac{x-a+1}{r} \rfloor & \text{if } a \leq x \leq b \\ 1 & \text{otherwise} \end{cases}$	(27.110)
mean	$EX = \frac{a+b}{2}$	(27.111)
median	$\text{med} = \frac{a+b}{2}$	(27.112)
mode	$\text{mode} = \emptyset$	(27.113)
variance	$D^2 X = \frac{r^2-1}{12}$	(27.114)
skewness	$\gamma_1 = 0$	(27.115)
kurtosis	$\gamma_2 = -\frac{6(r^2+1)}{5(r^2-1)}$	(27.116)
entropy	$H(X) = \ln r$	(27.117)
mgf	$M_X(t) = \frac{e^{at} - e^{(b+1)t}}{r(1-e^t)}$	(27.118)
char. func.	$\varphi_X(t) = \frac{e^{iat} - e^{i(b+1)t}}{r(1-e^{it})}$	(27.119)

Table 27.2: Parameters of the discrete uniform distribution.

⁴¹ http://en.wikipedia.org/wiki/Uniform_distribution_%28discrete%29 [accessed 2007-07-03]




Fig. 27.1: The PMFs of some discrete uniform distributions




Fig. 27.2: The CDFs of some discrete uniform distributions

27.3.2 Poisson Distribution π_λ

The Poisson distribution⁴² π_λ [1724] applies to the reference model *telephone switchboard*. It describes a process where the number of events that occur (independently of each other) in a certain time interval only depends on the duration of the interval and not of its position (*prehistory*). Events do not have any aftermath and thus, there is no mutual influence of non-overlapping time intervals (*homogeneity*). Furthermore, only the time when an even occurs is considered and not the duration of the event. In the telephone switchboard example, we would only be interested in the time at which a call comes in, not in the length of the call. In this model, no events occur in infinitely short time intervals. The features of the Poisson

⁴² http://en.wikipedia.org/wiki/Poisson_distribution [accessed 2007-07-03]

distribution are listed in Table 27.3⁴³ and examples for its PDF and CDF are illustrated in Figure 27.3 and Figure 27.4.

parameter	definition	
parameters	$\lambda = \mu t > 0$	(27.120)
PMF	$P(X = x) = f_X(x) = \frac{(\mu t)^x}{x!} e^{-\mu t} = \frac{\lambda^x}{x!} e^{-\lambda}$	(27.121)
CDF	$P(X \leq x) = F_X(x) = \frac{\Gamma(\lfloor k+1 \rfloor, \lambda)}{\lfloor k \rfloor!} = \sum_{i=0}^x \frac{e^{-\lambda} \lambda^i}{i!}$	(27.122)
mean	$EX = \mu t = \lambda$	(27.123)
median	$\text{med} \approx \lfloor \lambda + \frac{1}{3} - \frac{1}{5\lambda} \rfloor$	(27.124)
mode	$\text{mode} = \lfloor \lambda \rfloor$	(27.125)
variance	$D^2 X = \mu t = \lambda$	(27.126)
skewness	$\gamma_1 = \lambda^{-\frac{1}{2}}$	(27.127)
kurtosis	$\gamma_2 = \frac{1}{\lambda}$	(27.128)
entropy	$H(X) = \lambda(1 - \ln \lambda) + e^{-\lambda} \sum_{k=0}^{\infty} \frac{\lambda^k \ln(k!)}{k!}$	(27.129)
mgf	$M_X(t) = e^{\lambda(e^t - 1)}$	(27.130)
char. func.	$\varphi_X(t) = e^{\lambda(e^{it} - 1)}$	(27.131)

Table 27.3: Parameters of the Poisson distribution.




Fig. 27.3: The PMFs of some Poisson distributions

Poisson Process

The Poisson process⁴⁴ [1725] is a process that obeys the Poisson distribution – just like the example of the telephone switchboard mentioned before. Here, λ is expressed as the product of the intensity μ and the time t . μ normally describes a frequency, for example $\mu = \frac{1}{\text{min}}$.

⁴³ The Γ in Equation 27.122 denotes the (upper) incomplete gamma function. More information on the gamma function Γ can be found in Section 27.8.1 on page 466.

⁴⁴ http://en.wikipedia.org/wiki/Poisson_process [accessed 2007-07-03]




Fig. 27.4: The CDFs of some Poisson distributions

Both, the expected value as well as the variance of the Poisson process are $\lambda = \mu t$. In Equation 27.132, the probability that k events occur in a Poisson process in a time interval of the length t is defined.

$$P(X_t = k) = \frac{(\mu t)^k}{k!} e^{-\mu t} = \frac{\lambda^k}{k!} e^{-\lambda} \quad (27.132)$$

The probability that in a time interval $[t, t + \Delta t]$

- no events occur is $1 - \lambda \Delta t + o(\Delta t)$.
- exactly one event occurs is $\lambda \Delta t + o(\Delta t)$.
- multiple events occur $o(\Delta t)$.

Here we use an infinitesimal version the small- o notation.⁴⁵ The statement that $f \in o(\xi) \Rightarrow |f(x)| \ll |\xi(x)|$ is normally only valid for $x \rightarrow \infty$. In the infinitesimal variant, it holds for $x \rightarrow 0$. Thus, we can state that $o(\Delta t)$ is much smaller than Δt . In principle, the above equations imply that in an infinite small time span either no or one event occurs, i.e., events do not arrive simultaneously:

$$\lim_{t \rightarrow 0} P(X_t > 1) = 0 \quad (27.133)$$

The Relation between the Poisson Process and the Exponential Distribution

It is important to know that the (time) distance between two events of the Poisson process is exponentially distributed (see Section 27.4.3 on page 440). The expected value of the number of events to arrive per time unit in a Poisson process is EX_{pois} , then the expected value of the time between two events $\frac{1}{EX_{pois}}$. Since this is the expected value $EX_{exp} = \frac{1}{EX_{pois}}$ of the exponential distribution, its λ_{exp} -value is $\lambda_{exp} = \frac{1}{EX_{exp}} = \frac{1}{\frac{1}{EX_{pois}}} = EX_{pois}$. Therefore, the

⁴⁵ See Section 29.1.3 on page 482 and Definition 29.16 on page 483 for a detailed elaboration on the small- o notation.

λ_{exp} -value of the exponential distribution equals the λ_{pois} -value of the Poisson distribution $\lambda_{exp} = \lambda_{pois} = EX_{pois}$. In other words, the time interval between (neighboring) events of the Poisson process is exponentially distributed with the same λ value as the Poisson process, as illustrated in Equation 27.134.

$$X_i \sim \pi_\lambda \Leftrightarrow (t(X_{i+1}) - tX_i) \sim exp(\lambda) \quad \forall i \in \mathbb{N} \quad (27.134)$$

27.3.3 Binomial Distribution $B(n, p)$

The binomial distribution⁴⁶ $B(n, p)$ is the probability distribution that describes the probability of the possible numbers successes of n independent experiments with the success probability p each. Such experiments is called Bernoulli experiments or Bernoulli trials. For $n = 1$, the binomial distribution is a Bernoulli distribution⁴⁷.

Table 27.4⁴⁸ points out some of the properties of the binomial distribution. A few examples for PMFs and CDFs of different binomial distributions are given in Figure 27.5 and Figure 27.6.

parameter	definition	
parameters	$n \in \mathbb{N}_0, 0 \leq p \leq 1, p \in \mathbb{R}$	(27.135)
PMF	$P(X = x) = f_X(x) = \binom{n}{x} p^x (1-p)^{n-x}$	(27.136)
CDF	$P(X \leq x) = F_X(x) = \sum_{i=0}^{\lfloor x \rfloor} f_X(x) = I_{1-p}(n - \lfloor x \rfloor, 1 + \lfloor x \rfloor)$	(27.137)
mean	$EX = np$	(27.138)
median	med is one of $\{\lfloor np \rfloor - 1, \lfloor np \rfloor, \lfloor np \rfloor + 1\}$	(27.139)
mode	mode = $\lfloor (n+1)p \rfloor$	(27.140)
variance	$D^2 X = np(1-p)$	(27.141)
skewness	$\gamma_1 = \frac{1-2p}{\sqrt{np(1-p)}}$	(27.142)
kurtosis	$\gamma_2 = \frac{1-6p(1-p)}{np(1-p)}$	(27.143)
entropy	$H(X) = \frac{1}{2} \ln(2\pi nep(1-p)) + \mathcal{O}\left(\frac{1}{n}\right)$	(27.144)
mgf	$M_X(t) = (1-p + pe^t)^n$	(27.145)
char. func.	$\varphi_X(t) = (1-p + pe^{it})^n$	(27.146)

Table 27.4: Parameters of the Binomial distribution.

For $n \rightarrow \infty$, the binomial distribution approaches a normal distribution. For large n , $B(n, p)$ can therefore often be approximated with the normal distribution (see Section 27.4.2) $N(np, np(1-p))$. Whether this approximation is good or not can be found out by rules of thumb, some of them are:

$$\begin{aligned} np &> 5 \wedge n(1-p) > 5 \\ \mu \pm 3\sigma &\approx np \pm 3\sqrt{np(1-p)} \in [0, n] \end{aligned}$$

In case these rules hold, we still need to transform a continuous distribution to a discrete one. In order to do so, we add 0.5 to the x values, i. e., $F_{X,bin}(x) \approx F_{X,normal}(x + 0.5)$.

⁴⁶ http://en.wikipedia.org/wiki/Binomial_distribution [accessed 2007-10-01]

⁴⁷ http://en.wikipedia.org/wiki/Bernoulli_distribution [accessed 2007-10-01]

⁴⁸ I_{1-p} in Equation 27.137 denotes the regularized incomplete beta function.




Fig. 27.5: The PMFs of some binomial distributions




Fig. 27.6: The CDFs of some binomial distributions

27.4 Some Continuous Distributions

In this section we will introduce some common continuous distributions. Unlike the discrete distributions, continuous distributions have an uncountable infinite large set of possible outcomes of random experiments. Thus, the PDF does not assign probabilities to certain events. Only the CDF makes statements about the probability of a sub-set of possible outcomes of a random experiment.

27.4.1 Continuous Uniform Distribution

After discussing the discrete uniform distribution in Section 27.3.1, we now elaborate on its continuous form⁴⁹.

In a uniform distribution, all possible outcomes in a range $[a, b]$, $b > a$ have exactly the same probability. The characteristics of this distribution can be found in Table 27.5. Examples of its probability density function is illustrated in Figure 27.7 whereas the according cumulative density functions are outlined Figure 27.8.

parameter	definition	
parameters	$a, b \in \mathbb{R}, a \geq b$	(27.147)
PDF	$f_X(x) = \begin{cases} \frac{1}{b-a} & \text{if } x \in [a, b] \\ 0 & \text{otherwise} \end{cases}$	(27.148)
CDF	$P(X \leq x) = F_X(x) = \begin{cases} 0 & \text{if } x < a \\ \frac{x-a}{b-a} & \text{if } x \in [a, b] \\ 1 & \text{otherwise} \end{cases}$	(27.149)
mean	$EX = \frac{1}{2}(a + b)$	(27.150)
median	$\text{med} = \frac{1}{2}(a + b)$	(27.151)
mode	$\text{mode} = \emptyset$	(27.152)
variance	$D^2X = \frac{1}{12}(b - a)^2$	(27.153)
skewness	$\gamma_1 = 0$	(27.154)
kurtosis	$\gamma_2 = -\frac{6}{5}$	(27.155)
entropy	$h(X) = \ln(b - a)$	(27.156)
mgf	$M_X(t) = \frac{e^{tb} - e^{ta}}{t(b-a)}$	(27.157)
char. func.	$\varphi_X(t) = \frac{e^{itb} - e^{ita}}{it(b-a)}$	(27.158)

Table 27.5: Parameters of the continuous uniform distribution.




Fig. 27.7: The PDFs of some continuous uniform distributions

⁴⁹ http://en.wikipedia.org/wiki/Uniform_distribution_%28continuous%29 [accessed 2007-07-03]




Fig. 27.8: The CDFs of some continuous uniform distributions

27.4.2 Normal Distribution $N(\mu, \sigma^2)$

Many phenomena in nature, like the size of chicken eggs, noise, errors in measurement, and such and such, can be considered as outcomes of random experiments with properties that can be approximated by the normal distribution⁵⁰ $N(\mu, \sigma^2)$ [1726]. Its probability density function, shown for some example values in Figure 27.9, is symmetric to the expected value μ and becomes flatter with rising standard deviation σ . The cumulative density function is outline for the same example values in Figure 27.10. Other characteristics of the normal distribution can be found in Table 27.6.

parameter	definition	
parameters	$\mu \in \mathbb{R}, \sigma \in \mathbb{R}^+$	(27.159)
PDF	$f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$	(27.160)
CDF	$P(X \leq x) = F_X(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{(z-\mu)^2}{2\sigma^2}} dz$	(27.161)
mean	$EX = \mu$	(27.162)
median	$\text{med} = \mu$	(27.163)
mode	$\text{mode} = \mu$	(27.164)
variance	$D^2 X = \sigma^2$	(27.165)
skewness	$\gamma_1 = 0$	(27.166)
kurtosis	$\gamma_2 = 0$	(27.167)
entropy	$h(X) = \ln(\sigma\sqrt{2\pi e})$	(27.168)
mgf	$M_X(t) = e^{\mu t + \frac{\sigma^2 t^2}{2}}$	(27.169)
char. func.	$\varphi_X(t) = e^{\mu it + \frac{\sigma^2 t^2}{2}}$	(27.170)

Table 27.6: Parameters of the normal distribution.

Definition 27.48 (Standard Normal Distribution).

For the sake of simplicity, the standard normal distribution $N(0, 1)$ with the CDF $\Phi(x)$ is defined with $\mu = 0$ and $\sigma = 1$. Values of this function are listed in tables. You can compute

⁵⁰ http://en.wikipedia.org/wiki/Normal_distribution [accessed 2007-07-03]




Fig. 27.9: The PDFs of some normal distributions




Fig. 27.10: The CDFs of some normal distributions

the CDF of any normal distribution using the one of the standard normal distribution by applying Equation 27.171.

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{z^2}{2}} dz \quad (27.171)$$

$$P(X \leq x) = \Phi\left(\frac{x - \mu}{\sigma}\right) \quad (27.172)$$

Some values of $\Phi(x)$ are listed in Table 27.7. For the sake of saving space by using two dimensions, we compose the values of x as a sum of a row and column value. If you want to look up $\Phi(2.13)$ for example, you'd go to the row which starts with 2.1 and the column of 0.03, so you'd find $\Phi(2.13) \approx 0.9834$.

x	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0.0	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
0.1	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
0.2	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
0.3	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
0.6	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
0.7	0.7580	0.7611	0.7642	0.7673	0.7704	0.7734	0.7764	0.7794	0.7823	0.7852
0.8	0.7881	0.7910	0.7939	0.7967	0.7995	0.8023	0.8051	0.8078	0.8106	0.8133
0.9	0.8159	0.8186	0.8212	0.8238	0.8264	0.8289	0.8315	0.8340	0.8365	0.8389
1.0	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554	0.8577	0.8599	0.8621
1.1	0.8643	0.8665	0.8686	0.8708	0.8729	0.8749	0.8770	0.8790	0.8810	0.8830
1.2	0.8849	0.8869	0.8888	0.8907	0.8925	0.8944	0.8962	0.8980	0.8997	0.9015
1.3	0.9032	0.9049	0.9066	0.9082	0.9099	0.9115	0.9131	0.9147	0.9162	0.9177
1.4	0.9192	0.9207	0.9222	0.9236	0.9251	0.9265	0.9279	0.9292	0.9306	0.9319
1.5	0.9332	0.9345	0.9357	0.9370	0.9382	0.9394	0.9406	0.9418	0.9429	0.9441
1.6	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515	0.9525	0.9535	0.9545
1.7	0.9554	0.9564	0.9573	0.9582	0.9591	0.9599	0.9608	0.9616	0.9625	0.9633
1.8	0.9641	0.9649	0.9656	0.9664	0.9671	0.9678	0.9686	0.9693	0.9699	0.9706
1.9	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750	0.9756	0.9761	0.9767
2.0	0.9772	0.9778	0.9783	0.9788	0.9793	0.9798	0.9803	0.9808	0.9812	0.9817
2.1	0.9821	0.9826	0.9830	0.9834	0.9838	0.9842	0.9846	0.9850	0.9854	0.9857
2.2	0.9861	0.9864	0.9868	0.9871	0.9875	0.9878	0.9881	0.9884	0.9887	0.9890
2.3	0.9893	0.9896	0.9898	0.9901	0.9904	0.9906	0.9909	0.9911	0.9913	0.9916
2.4	0.9918	0.9920	0.9922	0.9925	0.9927	0.9929	0.9931	0.9932	0.9934	0.9936
2.5	0.9938	0.9940	0.9941	0.9943	0.9945	0.9946	0.9948	0.9949	0.9951	0.9952
2.6	0.9953	0.9955	0.9956	0.9957	0.9959	0.9960	0.9961	0.9962	0.9963	0.9964
2.7	0.9965	0.9966	0.9967	0.9968	0.9969	0.9970	0.9971	0.9972	0.9973	0.9974
2.8	0.9974	0.9975	0.9976	0.9977	0.9977	0.9978	0.9979	0.9979	0.9980	0.9981
2.9	0.9981	0.9982	0.9982	0.9983	0.9984	0.9984	0.9985	0.9985	0.9986	0.9986
3.0	0.9987	0.9987	0.9987	0.9988	0.9988	0.9989	0.9989	0.9989	0.9990	0.9990

Table 27.7: Some values of the standardized normal distribution.

Definition 27.49 (probit). The inverse of the cumulative distribution function of the standard normal distribution is called the probit function. It is also often denoted as z -quantile of the standard normal distribution.

$$z(y) \equiv \text{probit}(y) \equiv \Phi^{-1}(y) \quad (27.173)$$

$$y = \Phi(x) \Rightarrow \Phi^{-1}(y) = z(y) = x \quad (27.174)$$

The probability density function PDF of the multivariate normal distribution⁵¹ [1727, 1728, 1729] is illustrated in Equation 27.175 and Equation 27.176 in the general case (where Σ is the covariance matrix) and in Equation 27.177 in the uncorrelated form. If the distributions, additionally to being uncorrelated, also have the same parameters σ and μ , the probability density function of the multivariate normal distribution can be expressed as it is done in Equation 27.178.

⁵¹ http://en.wikipedia.org/wiki/Multivariate_normal_distribution [accessed 2007-07-03]

$$f_X(\mathbf{x}) = \frac{\sqrt{\Sigma^{-1}}}{(2\pi)^{\frac{n}{2}}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{x}-\boldsymbol{\mu})} \quad (27.175)$$

$$= \frac{1}{(2\pi)^{\frac{n}{2}} \Sigma^{\frac{1}{2}}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{x}-\boldsymbol{\mu})} \quad (27.176)$$

$$f_X(\mathbf{x}) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi}\sigma_i} e^{-\frac{(x_i-\mu_i)^2}{2\sigma_i^2}} \quad (27.177)$$

$$\begin{aligned} f_X(\mathbf{x}) &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x_i-\mu_i)^2}{2\sigma^2}} \\ &= \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{n}{2}} e^{-\frac{\sum_{i=1}^n (x_i-\mu)^2}{2\sigma^2}} \end{aligned} \quad (27.178)$$

27.4.3 Exponential Distribution $\exp(\lambda)$

The exponential distribution⁵² $\exp(\lambda)$ [1730] is often used if the probabilities of lifetimes of apparatuses, half-life periods of radioactive elements, or the time between two events in the Poisson process (see Section 27.3.2 on page 433) has to be approximated. Its PDF is sketched in Figure 27.11 for some example values of λ the according cases of the CDF are illustrated Figure 27.12. The most important characteristics of the exponential distribution can be obtained from Table 27.8.

parameter	definition	
parameters	$\lambda \in \mathbb{R}^+$	(27.179)
PDF	$f_X(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ \lambda e^{-\lambda x} & \text{otherwise} \end{cases}$	(27.180)
CDF	$P(X \leq x) = F_X(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ 1 - e^{-\lambda x} & \text{otherwise} \end{cases}$	(27.181)
mean	$EX = \frac{1}{\lambda}$	(27.182)
median	$\text{med} = \frac{\ln 2}{\lambda}$	(27.183)
mode	$\text{mode} = 0$	(27.184)
variance	$D^2 X = \frac{1}{\lambda^2}$	(27.185)
skewness	$\gamma_1 = 2$	(27.186)
kurtosis	$\gamma_2 = 6$	(27.187)
entropy	$h(X) = 1 - \ln \lambda$	(27.188)
mgf	$M_X(t) = \left(1 - \frac{t}{\lambda}\right)^{-1}$	(27.189)
char. func.	$\varphi_X(t) = \left(1 - \frac{it}{\lambda}\right)^{-1}$	(27.190)

Table 27.8: Parameters of the exponential distribution.

⁵² http://en.wikipedia.org/wiki/Exponential_distribution [accessed 2007-07-03]




Fig. 27.11: The PDFs of some exponential distributions




Fig. 27.12: The CDFs of some exponential distributions

27.4.4 Chi-square Distribution

The chi-square (or χ^2) distribution⁵³ is a steady probability distribution on the set of positive real numbers. It is a so-called *sample distribution* which is used for the estimation of parameters like the variance of other distributions. We can also describe the sum of independent standardized normal distributions with it. Its sole parameter, n , denotes the degrees of freedom.

In Table 27.9⁵⁴, the characteristic parameters of the χ^2 distribution are outlined. A few examples for the PDF and CDF of the χ^2 distribution are illustrated in Figure 27.13 and Figure 27.14.

⁵³ http://en.wikipedia.org/wiki/Chi-square_distribution [accessed 2007-09-30]

⁵⁴ $\gamma(n, z)$ in Equation 27.193 is the lower incomplete Gamma function and $P_\gamma(n, z)$ is the regularized Gamma function.

Table 27.10 provides some selected values of the χ^2 distribution. The table's headline contains results of the cumulative distribution function $F_X(x)$ of a χ^2 distribution with n degrees of freedom (values in the first column). The cells now denote the x values that belong to these $(m, F_X(x))$ combinations.

parameter	definition	
parameters	$n \in \mathbb{R}^+, n > 0$	(27.191)
PDF	$f_X(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ \frac{2^{-n/2}}{\Gamma(n/2)} x^{n/2-1} e^{-x/2} & \text{otherwise} \end{cases}$	(27.192)
CDF	$P(X \leq x) = F_X(x) = \frac{\gamma(n/2, x/2)}{\Gamma(n/2)} = P_\gamma(n/2, x/2)$	(27.193)
mean	$EX = n$	(27.194)
median	$\text{med} \approx n - \frac{2}{3}$	(27.195)
mode	$\text{mode} = n - 2 \text{ if } n \geq 2$	(27.196)
variance	$D^2 X = 2n$	(27.197)
skewness	$\gamma_1 = \sqrt{\frac{8}{n}}$	(27.198)
kurtosis	$\gamma_2 = \frac{12}{n}$	(27.199)
entropy	$h(X) = \frac{n}{2} + \ln(2\Gamma(n/2)) + (1 - n/2)\psi(n/2)$	(27.200)
mgf	$M_X(t) = (1 - 2t)^{-n/2} \text{ for } 2t < 1$	(27.201)
char. func.	$\varphi_X(t) = (1 - 2it)^{-n/2}$	(27.202)

Table 27.9: Parameters of the χ^2 distribution.





Fig. 27.13: The PDFs of some χ^2 distributions

Fig. 27.14: The CDFs of some χ^2 distributions

n	0.995	.99	.975	.95	.9	.1	.05	.025	.01	.005
1	–	–	0.001	0.004	0.016	2.706	3.841	5.024	6.635	7.879
2	0.010	0.020	0.051	0.103	0.211	4.605	5.991	7.378	9.210	10.597
3	0.072	0.115	0.216	0.352	0.584	6.251	7.815	9.348	11.345	12.838
4	0.207	0.297	0.484	0.711	1.064	7.779	9.488	11.143	13.277	14.860
5	0.412	0.554	0.831	1.145	1.610	9.236	11.070	12.833	15.086	16.750
6	0.676	0.872	1.237	1.635	2.204	10.645	12.592	14.449	16.812	18.548
7	0.989	1.239	1.690	2.167	2.833	12.017	14.067	16.013	18.475	20.278
8	1.344	1.646	2.180	2.733	3.490	13.362	15.507	17.535	20.090	21.955
9	1.735	2.088	2.700	3.325	4.168	14.684	16.919	19.023	21.666	23.589
10	2.156	2.558	3.247	3.940	4.865	15.987	18.307	20.483	23.209	25.188
11	2.603	3.053	3.816	4.575	5.578	17.275	19.675	21.920	24.725	26.757
12	3.074	3.571	4.404	5.226	6.304	18.549	21.026	23.337	26.217	28.300
13	3.565	4.107	5.009	5.892	7.042	19.812	22.362	24.736	27.688	29.819
14	4.075	4.660	5.629	6.571	7.790	21.064	23.685	26.119	29.141	31.319
15	4.601	5.229	6.262	7.261	8.547	22.307	24.996	27.488	30.578	32.801
16	5.142	5.812	6.908	7.962	9.312	23.542	26.296	28.845	32.000	34.267
17	5.697	6.408	7.564	8.672	10.085	24.769	27.587	30.191	33.409	35.718
18	6.265	7.015	8.231	9.390	10.865	25.989	28.869	31.526	34.805	37.156
19	6.844	7.633	8.907	10.117	11.651	27.204	30.144	32.852	36.191	38.582
20	7.434	8.260	9.591	10.851	12.443	28.412	31.410	34.170	37.566	39.997
21	8.034	8.897	10.283	11.591	13.240	29.615	32.671	35.479	38.932	41.401
22	8.643	9.542	10.982	12.338	14.041	30.813	33.924	36.781	40.289	42.796
23	9.260	10.196	11.689	13.091	14.848	32.007	35.172	38.076	41.638	44.181
24	9.886	10.856	12.401	13.848	15.659	33.196	36.415	39.364	42.980	45.559
25	10.520	11.524	13.120	14.611	16.473	34.382	37.652	40.646	44.314	46.928
26	11.160	12.198	13.844	15.379	17.292	35.563	38.885	41.923	45.642	48.290
27	11.808	12.879	14.573	16.151	18.114	36.741	40.113	43.195	46.963	49.645
28	12.461	13.565	15.308	16.928	18.939	37.916	41.337	44.461	48.278	50.993
29	13.121	14.256	16.047	17.708	19.768	39.087	42.557	45.722	49.588	52.336
30	13.787	14.953	16.791	18.493	20.599	40.256	43.773	46.979	50.892	53.672
40	20.707	22.164	24.433	26.509	29.051	51.805	55.758	59.342	63.691	66.766
50	27.991	29.707	32.357	34.764	37.689	63.167	67.505	71.420	76.154	79.490
60	35.534	37.485	40.482	43.188	46.459	74.397	79.082	83.298	88.379	91.952
70	43.275	45.442	48.758	51.739	55.329	85.527	90.531	95.023	100.425	104.215
80	51.172	53.540	57.153	60.391	64.278	96.578	101.879	106.629	112.329	116.321
90	59.196	61.754	65.647	69.126	73.291	107.565	113.145	118.136	124.116	128.299
100	67.328	70.065	74.222	77.929	82.358	118.498	124.342	129.561	135.807	140.169

Table 27.10: Some values of the χ^2 distribution.

27.4.5 Student's t-Distribution

The Student's t-distribution⁵⁵ is based on the insight that the mean of a normally distributed feature of a sample is no longer normally distributed if the variance is unknown and needs to be estimated from the data samples [1731, 1732, 1733]. It has been design by Gosset [1731] who published it under the pseudonym *Student*.

The parameter n of the distribution denotes the degrees of freedom of the distribution. If n approaches infinity, the t-distribution approaches the standard normal distribution.

The characteristic properties of Student's t-distribution are outlined in Table 27.11⁵⁶ and examples for its PDF and CDF are illustrated in Figure 27.15 and Figure 27.16.

Table 27.12 provides some selected values for the quantiles $t_{1-\alpha,n}$ of the t-distribution (one-sided confidence intervals, see Section 27.6.3 on page 454). The headline of the table contains results of the cumulative distribution function $F_X(x)$ of a Student's t-distribution with n degrees of freedom (values in the first column). The cells now denote the x values that belong to these $(n, F_X(x))$ combinations.

parameter	definition	
parameters	$n \in \mathbb{R}^+, n > 0$	(27.203)
PDF	$f_X(x) = \frac{\Gamma((n+1)/2)}{\sqrt{n\pi}\Gamma(n/2)} (1 + x^2/n)^{-(n+1)/2}$	(27.204)
CDF	$P(X \leq x) = F_X(x) = \frac{1}{2} + x\Gamma\left(\frac{n+1}{2}\right) \frac{{}_2F_1\left(\frac{1}{2}, \frac{n+1}{2}, \frac{3}{2}, -\frac{x^2}{n}\right)}{\sqrt{n\pi}\Gamma\left(\frac{n}{2}\right)}$	(27.205)
mean	$EX = 0$	(27.206)
median	$\text{med} = 0$	(27.207)
mode	$\text{mode} = 0$	(27.208)
variance	$D^2X = \frac{n}{n-2}$ for $n > 2$, otherwise undefined	(27.209)
skewness	$\gamma_1 = 0$ for $n > 3$	(27.210)
kurtosis	$\gamma_2 = \frac{6}{n-4}$ for $n > 4$	(27.211)
entropy	$h(X) = \frac{n}{2} [\psi(\frac{n+1}{2}) - \psi(\frac{n}{2})] + \log [\sqrt{n}B(\frac{n}{2}, \frac{1}{2})]$	(27.212)
mgf	undefined	(27.213)

Table 27.11: Parameters of the Student's t- distribution.

⁵⁵ http://en.wikipedia.org/wiki/Student%27s_t-distribution [accessed 2007-09-30]

⁵⁶ More information on the gamma function Γ used in Equation 27.204 and Equation 27.205 can be found in Section 27.8.1 on page 466. ${}_2F_1$ in Equation 27.205 stands for the hypergeometric function, ψ and B in Equation 27.212 are the digamma and the beta function.




Fig. 27.15: The PDFs of some Student's t-distributions




Fig. 27.16: The CDFs of some Student's t-distributions

n	0.75	.8	.85	.875	.9	.95	.975	.99	.995	.9975	.999	.9995
1	1.000	1.376	1.963	2.414	3.078	6.314	12.71	31.82	63.66	127.3	318.3	636.6
2	0.816	1.061	1.386	1.605	1.886	2.920	4.303	6.965	9.925	14.09	22.33	31.60
3	0.765	0.978	1.250	1.423	1.638	2.353	3.182	4.541	5.841	7.453	10.21	12.92
4	0.741	0.941	1.190	1.344	1.533	2.132	2.776	3.747	4.604	5.598	7.173	8.610
5	0.727	0.920	1.156	1.301	1.476	2.015	2.571	3.365	4.032	4.773	5.893	6.869
6	0.718	0.906	1.134	1.273	1.440	1.943	2.447	3.143	3.707	4.317	5.208	5.959
7	0.711	0.896	1.119	1.254	1.415	1.895	2.365	2.998	3.499	4.029	4.785	5.408
8	0.706	0.889	1.108	1.240	1.397	1.860	2.306	2.896	3.355	3.833	4.501	5.041
9	0.703	0.883	1.100	1.230	1.383	1.833	2.262	2.821	3.250	3.690	4.297	4.781
10	0.700	0.879	1.093	1.221	1.372	1.812	2.228	2.764	3.169	3.581	4.144	4.587
11	0.697	0.876	1.088	1.214	1.363	1.796	2.201	2.718	3.106	3.497	4.025	4.437
12	0.695	0.873	1.083	1.209	1.356	1.782	2.179	2.681	3.055	3.428	3.930	4.318
13	0.694	0.870	1.079	1.204	1.350	1.771	2.160	2.650	3.012	3.372	3.852	4.221
14	0.692	0.868	1.076	1.200	1.345	1.761	2.145	2.624	2.977	3.326	3.787	4.140
15	0.691	0.866	1.074	1.197	1.341	1.753	2.131	2.602	2.947	3.286	3.733	4.073
16	0.690	0.865	1.071	1.194	1.337	1.746	2.120	2.583	2.921	3.252	3.686	4.015
17	0.689	0.863	1.069	1.191	1.333	1.740	2.110	2.567	2.898	3.222	3.646	3.965
18	0.688	0.862	1.067	1.189	1.330	1.734	2.101	2.552	2.878	3.197	3.610	3.922
19	0.688	0.861	1.066	1.187	1.328	1.729	2.093	2.539	2.861	3.174	3.579	3.883
20	0.687	0.860	1.064	1.185	1.325	1.725	2.086	2.528	2.845	3.153	3.552	3.850
21	0.686	0.859	1.063	1.183	1.323	1.721	2.080	2.518	2.831	3.135	3.527	3.819
22	0.686	0.858	1.061	1.182	1.321	1.717	2.074	2.508	2.819	3.119	3.505	3.792
23	0.685	0.858	1.060	1.180	1.319	1.714	2.069	2.500	2.807	3.104	3.485	3.767
24	0.685	0.857	1.059	1.179	1.318	1.711	2.064	2.492	2.797	3.091	3.467	3.745
25	0.684	0.856	1.058	1.178	1.316	1.708	2.060	2.485	2.787	3.078	3.450	3.725
26	0.684	0.856	1.058	1.177	1.315	1.706	2.056	2.479	2.779	3.067	3.435	3.707
27	0.684	0.855	1.057	1.176	1.314	1.703	2.052	2.473	2.771	3.057	3.421	3.690
28	0.683	0.855	1.056	1.175	1.313	1.701	2.048	2.467	2.763	3.047	3.408	3.674
29	0.683	0.854	1.055	1.174	1.311	1.699	2.045	2.462	2.756	3.038	3.396	3.659
30	0.683	0.854	1.055	1.173	1.310	1.697	2.042	2.457	2.750	3.030	3.385	3.646
40	0.681	0.851	1.050	1.167	1.303	1.684	2.021	2.423	2.704	2.971	3.307	3.551
50	0.679	0.849	1.047	1.164	1.299	1.676	2.009	2.403	2.678	2.937	3.261	3.496
60	0.679	0.848	1.045	1.162	1.296	1.671	2.000	2.390	2.660	2.915	3.232	3.460
80	0.678	0.846	1.043	1.159	1.292	1.664	1.990	2.374	2.639	2.887	3.195	3.416
100	0.677	0.845	1.042	1.158	1.290	1.660	1.984	2.364	2.626	2.871	3.174	3.390
120	0.677	0.845	1.041	1.157	1.289	1.658	1.980	2.358	2.617	2.860	3.160	3.373
∞	0.674	0.842	1.036	1.150	1.282	1.645	1.960	2.326	2.576	2.807	3.090	3.291

Table 27.12: Table of Student's t-distribution with right-tail probabilities.

27.5 Example – Throwing a Dice

Let us now discuss the different parameters of a random variable at the example of throwing a dice. On a dice, numbers from one to six are written and the result of throwing it is the number written on the side facing upwards. If a dice is perfect, the numbers one to six will show up with exactly the same probability, $\frac{1}{6}$. The set of all possible outcomes of throwing a dice Ω is thus

$$\Omega = \{\boxed{1}, \boxed{2}, \boxed{3}, \boxed{4}, \boxed{5}, \boxed{6}\} \quad (27.214)$$

We define a random variable $X : \Omega \mapsto \mathbb{R}$ that assigns real numbers to the possible outcomes of throwing the dice in a way that the value of X matches the number on the dice:

$$X : \Omega \mapsto \{1, 2, 3, 4, 5, 6\} \quad (27.215)$$

It is obviously a uniformly distributed discrete random variable (see Section 27.3.1 on page 430) that can take on six states. We can now define the probability mass function PMF and the according cumulative distribution function CDF as follows (see also Figure 27.17):

$$F_X(x) = P(X \leq x) = \begin{cases} 0 & \text{if } x < 1 \\ \frac{x}{6} & \text{if } 1 \leq x \leq 6 \\ 1 & \text{otherwise} \end{cases} \quad (27.216)$$

$$f_X(x) = P(X = x) = \begin{cases} 0 & \text{if } x < 1 \\ \frac{1}{6} & \text{if } 1 \leq x \leq 6 \\ 0 & \text{otherwise} \end{cases} \quad (27.217)$$




Fig. 27.17: The PMF and CMF of the dice throw

We now can discuss the statistical parameters of this experiment. This is a good opportunity to compare the real parameters and their estimates. We therefore assume that the dice was thrown ten times ($n = 10$) in an experiment. The following numbers have been thrown as illustrated in Figure 27.18):

$$A = \{4, 5, 3, 2, 4, 6, 4, 2, 5, 3\} \quad (27.218)$$

Table 27.13 outlines how the parameters of the random variable are computed. The real




Fig. 27.18: The numbers thrown in the dice example

parameter true value		estimate	
count	non existent	$n = \text{len}(A) = 10$	(27.219)
minimum	$a = \min \{x : f_X(x) > 0\} = 1$	$\tilde{a} = \min A = 2 \approx a$	(27.220)
maximum	$b = \max \{x : f_X(x) > 0\} = 6$	$\tilde{b} = \max A = 6 \approx b$	(27.221)
range	$\text{range} = r = b - a + 1 = 6$	$\tilde{r} = b - a + 1 = 6 \approx \text{range}$	(27.222)
mean	$EX = \frac{a+b}{2} = \frac{7}{2} = 3.5$	$\bar{a} = \frac{1}{n} \sum_{i=0}^{n-1} A[i] = \frac{19}{5} = 3.8 \approx EX$	(27.223)
median	$\text{med} = \frac{a+b}{2} = \frac{7}{2} = 3.5$	$A_s = \text{sortList}_a(A, >)$	(27.224)
		$\widehat{\text{med}} = \frac{A_s[\frac{n}{2}] + A_s[\frac{n}{2}-1]}{2} = 4 \approx \text{med}$	
mode	$\text{mode} = \emptyset$	$\text{mode} = \{4\} \approx \text{mode}$	(27.225)
variance	$D^2X = \frac{r^2-1}{12} = \frac{35}{12} \approx 2.917$	$s^2 = \frac{1}{n-1} \sum_{i=0}^{n-1} (A[i] - \bar{a})^2 = \frac{26}{15} \approx 1.73 \approx D^2X$	(27.226)
skewness	$\gamma_1 = 0$	$G_1 \approx 0.0876 \approx \gamma_1$	(27.227)
kurtosis	$\gamma_2 = -\frac{6(r^2+1)}{5(r^2-1)} = -\frac{222}{175} \approx -1.269$	$G_2 \approx -0.7512 \approx \gamma_2$	(27.228)

Table 27.13: Parameters of the dice throw experiment.

values of the parameters are defined using the PMF or CDF functions, while the estimations are based on the sample data obtained from our experiment solely.

As you can see, the estimations of the parameters sometimes differ significantly from their true values. More information about estimation can be found in the following section.

27.6 Estimation Theory

Estimation theory is the science of approximating the values of parameters based on measurements or otherwise obtained sample data [1734, 1735, 1736, 1737, 1738]. The center of this branch of statistics is to find good estimators in order to approximate the real values the parameters as good as possible.

Definition 27.50 (Estimator). An estimator⁵⁷ $\tilde{\theta}$ is a rule (most often a mathematical function) that takes a set of sample data A as input and returns an estimation of one parameter θ of the random distribution of the process sampled with this data set.

We have already discussed some estimators in Section 27.2 – the arithmetic mean of a sample data set (see Definition 27.29 on page 424) for example is an estimator for the expected value (see Definition 27.27 on page 424) and in Equation 27.67 on page 425 we have introduced an estimator for the sample variance.

Obviously, the estimator $\tilde{\theta}$ is the better the closer its results (the estimates) come to the real values of the parameter θ .

Definition 27.51 (Point Estimator). We define a point estimator $\tilde{\theta}$ to be an estimator which is a mathematical function $\tilde{\theta} : \mathbb{R}^n \mapsto \mathbb{R}$. This function takes the data sample A (here considered as a real vector $A \in \mathbb{R}^n$) as input and returns the estimate in the form of a (real) scalar value.

Definition 27.52 (Error). The absolute (estimation) error ε ⁵⁸ is the difference between the value returned by a point estimator $\tilde{\theta}$ of a parameter θ for a certain input A and its real value. Notice that the error ε can be zero, positive, or negative.

$$\varepsilon_A(\tilde{\theta}) = \tilde{\theta}(A) - \theta \quad (27.229)$$

In the following, we will most often not explicitly refer to the data sample A as basis of the estimation $\tilde{\theta}$ anymore. We assume that it is implicitly clear that estimations are usually based on such samples and that subscripts like the A in ε_A in Equation 27.229 are not needed.

Definition 27.53 (Bias). The bias $\text{Bias}(\tilde{\theta})$ of an estimator $\tilde{\theta}$ is the expected value of the difference of the estimate and the real value. This mean error is null for all unbiased estimators.

$$\text{Bias}(\tilde{\theta}) = E[\tilde{\theta} - \theta] = E[\varepsilon(\tilde{\theta})] \quad (27.230)$$

Definition 27.54 (Unbiased Estimator). An unbiased estimator has a zero bias.

$$\text{Bias}(\tilde{\theta}) = E[\tilde{\theta} - \theta] = E[\varepsilon(\tilde{\theta})] = 0 \Leftrightarrow E\tilde{\theta} = \theta \quad (27.231)$$

Definition 27.55 (Mean Square Error). The mean square error⁵⁹ $\text{MSE}(\tilde{\theta})$ of an estimator $\tilde{\theta}$ is the expected value of the square of the estimation error ε . It is also the sum of the variance of the estimator and the square of its bias. The MSE is a measure for how much an estimator differs from the quantity to be estimated.

$$\text{MSE}(\tilde{\theta}) = E[(\tilde{\theta} - \theta)^2] = E[(\varepsilon(\tilde{\theta}))^2] \quad (27.232)$$

$$\text{MSE}(\tilde{\theta}) = D^2\tilde{\theta} + (\text{Bias}(\tilde{\theta}))^2 \quad (27.233)$$

⁵⁷ <http://en.wikipedia.org/wiki/Estimator> [accessed 2007-07-03], <http://mathworld.wolfram.com/Estimator.html> [accessed 2007-07-03]

⁵⁸ http://en.wikipedia.org/wiki/Errors_and_residuals_in_statistics [accessed 2007-07-03]

⁵⁹ http://en.wikipedia.org/wiki/Mean_squared_error [accessed 2007-07-03]

Notice that the MSE of unbiased estimators coincides with the variance $D^2\tilde{\theta}$ of $\tilde{\theta}$. For estimating the mean square error of an estimator $\tilde{\theta}$, we use the sample mean:

$$\widehat{\text{MSE}}(\tilde{\theta}) = \frac{1}{n} \sum_{i=1}^n (\tilde{\theta}_i - \tilde{\theta})^2 \quad (27.234)$$

27.6.1 Likelihood and Maximum Likelihood Estimators

Definition 27.56 (Likelihood). Likelihood⁶⁰ is a mathematical expression complementary to probability. Whereas probability allows us to predict the outcome of a random experiment based on known parameters, likelihood allows us to predict unknown parameters based on the outcome of experiments.

Definition 27.57 (Likelihood Function). The likelihood function L returns a value that is proportional to the probability of a postulated underlying law or probability distribution φ according to an observed outcome (denoted as the vector $\mathbf{y}$). Notice that L not necessarily represents a probability density/mass function and its integral also does not necessarily equal to 1.

$$L[\varphi|\mathbf{y}] \propto P(\mathbf{y}|\varphi) \quad (27.235)$$

In many sources, L is defined in dependency of a parameter θ instead of the function φ . We preferred the latter notation since it is a more general superset of the first one.

Observation of an Unknown Process φ

Assume that we are given a finite set A of n sample data points.

$$A = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}, \quad x_i, y_i \in \mathbb{R} \quad \forall i \in [1, n] \quad (27.236)$$

The x_i are known inputs or parameters of an unknown process defined by the function $\varphi: \mathbb{R} \mapsto \mathbb{R}$. By observing the corresponding outputs of the process, we have obtained the y_i values. During our observations, we make the measurement errors⁶¹ η_i .

$$y_i = \varphi(x_i) + \eta_i \quad \forall i: 0 < i \leq n \quad (27.237)$$

About this measurement error η we make the following assumptions:

$$E\eta = 0 \quad (27.238)$$

$$\eta \sim N(0, \sigma^2) : 0 < \sigma < \infty \quad (27.239)$$

$$\text{cov}(\eta_i, \eta_j) = 0 \quad \forall i, j \in \mathbb{N} : i \neq j, 0 < i \leq n, 0 < j \leq n \quad (27.240)$$

- The expected values of η in Equation 27.238 are all zero. Our measurement device thus gives us, in average, unbiased results. If the expected value of η was not zero, we could simply recalibrate our (imaginary) measurement equipment in order to subtract $E\eta$ from all measurements and would obtain unbiased observations.
- Furthermore, Equation 27.239 states that the η_i are normally distributed around the zero point with an unknown, nonzero variance σ^2 . To suppose measurement errors to be normally distributed is quite common and correct in most cases. The white noise⁶² in transmission of signals for example is often modeled with Gaussian distributed⁶³ amplitudes. This second assumption includes, of course, the first one: Being normally distributed with $N(\mu = 0, \sigma^2)$ implies a zero expected value of the error.

⁶⁰ <http://en.wikipedia.org/wiki/Likelihood> [accessed 2007-07-03]

⁶¹ http://en.wikipedia.org/wiki/Measurement_error [accessed 2007-07-03]

⁶² http://en.wikipedia.org/wiki/White_noise [accessed 2007-07-03]

⁶³ http://en.wikipedia.org/wiki/Gaussian_noise [accessed 2007-07-03]

- With Equation 27.240, we assume that the errors η_i of the single measurements are stochastically independent. If there existed a connection between them, it would be part of the underlying physical law φ and could be incorporated in our measurement device and again be subtracted.

Objective: Estimation

Assume that we can choose from a, possibly infinite large, set of functions (estimators) $f \in F$.

$$f \in F \Rightarrow f : \mathbb{R} \mapsto \mathbb{R} \quad (27.241)$$

From this set we want to pick the function $f^* \in F$ with that resembles φ the best (i. e., better than all other $f \in F : f \neq f^*$). φ is not necessarily an element of F , so we cannot always presume to find a $f^* \equiv \varphi$.

Each estimator f deviates by the estimation error $\varepsilon(f)$ (see Definition 27.52 on page 450) from the y_i -values. The estimation error depends on f and may vary for different estimators.

$$y_i = f(x_i) + \varepsilon_i(f) \quad \forall i : 0 < i \leq n \quad (27.242)$$

We consider all $f \in F$ to be valid estimators for φ and simply look for the one that “fits best”. We now can combine Equation 27.242 with Equation 27.237:

$$f(x_i) + \varepsilon_i(f) = y_i = \varphi(x_i) + \eta_i \quad \forall i : 0 < i \leq n \quad (27.243)$$

We do not know φ and thus, cannot determine the η_i . According to the likelihood method, we pick the function $f \in F$ that would have most probably produced the outcomes y_i . In other words, we have to maximize the likelihood of the occurrence of the $\varepsilon_i(f)$. The likelihood here is defined under the assumption that the true measurement errors η_i are normally distributed (see Equation 27.239). So what we can do is to determine the ε_i in a way that their occurrence is most probable according to the distribution of the random variable that created the η_i , $N(0, \sigma^2)$. In the best case, the $\varepsilon(f^*) = \eta_i$ and thus, f^* is equivalent to $\varphi(x_i)$, at least in for the sample information A available to us.

Maximizing the Likelihood

Therefore, we can regard the $\varepsilon_i(f)$ as outcomes of independent random experiments, as uncorrelated random variables, and combine them to a multivariate normal distribution. For the ease of notation, we define the $\varepsilon(f)$ to be the vector containing all the single $\varepsilon_i(f)$ -values.

$$\varepsilon(f) = \begin{pmatrix} \varepsilon_1(f) \\ \varepsilon_2(f) \\ \vdots \\ \varepsilon_n(f) \end{pmatrix} \quad (27.244)$$

The probability density function of a multivariate normal distribution with independent variables ε_i that have the same variance σ^2 looks like this (as defined in Equation 27.178 on page 440):

$$f_X(\varepsilon(f)) = \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{n}{2}} e^{-\frac{\sum_{i=1}^n (\varepsilon_i(f) - \mu)^2}{2\sigma^2}} = \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{n}{2}} e^{-\frac{\sum_{i=1}^n (\varepsilon_i(f))^2}{2\sigma^2}} \quad (27.245)$$

Amongst all possible vectors $\varepsilon(f) : f \in F$ we need to find the most probable one $\varepsilon^* = \varepsilon(f^*)^*$ according to Equation 27.245. The function f^* that produces it will then be the one that most probably matches to φ .

In order to express how likely the observation of some outcomes is under a certain set of parameters, we have defined the likelihood function L in Definition 27.57. Here we can use the probability density function f_X of the normal distribution, since the maximal values of f_X are those that are most probable to occur.

$$L[\varepsilon(f)|f] = f_X(\varepsilon(f)) = \left(\frac{1}{2\pi\sigma^2}\right)^{\frac{n}{2}} e^{-\frac{\sum_{i=1}^n (\varepsilon_i(f))^2}{2\sigma^2}} \quad (27.246)$$

$$f^* \in F : L[\varepsilon(f^*)|f^*] = \max_{\forall f \in F} L[\varepsilon(f)|f] \quad (27.247)$$

$$= \max_{\forall f \in F} \left(\frac{1}{2\pi\sigma^2}\right)^{\frac{n}{2}} e^{-\frac{\sum_{i=1}^n (\varepsilon_i(f))^2}{2\sigma^2}} \quad (27.248)$$

Finding a f^* that Maximizes the function f_X however is equal to find a f^* that minimizes the sum of the squares of the ε -values.

$$f^* \in F : \sum_{i=1}^n (\varepsilon_i(f^*))^2 = \min_{\forall f \in F} \sum_{i=1}^n (\varepsilon_i(f))^2 \quad (27.249)$$

According to Equation 27.242 we can now substitute the ε_i -values with the difference between the observed outcomes y_i and the estimates $f(x_i)$.

$$\sum_{i=1}^n (\varepsilon_i(f))^2 = \sum_{i=1}^n (y_i - f(x_i))^2 \quad (27.250)$$

Definition 27.58 (Maximum Likelihood Estimator). A maximum likelihood estimator⁶⁴ [1739] f^* is an estimator which fits with maximum likelihood to a given set of sample data A . Under the particular assumption of uncorrelated error terms normally distributed around zero, a MLE minimizes Equation 27.251.

$$f^* \in F : \sum_{i=1}^n (y_i - f^*(x_i))^2 = \min_{\forall f \in F} \sum_{i=1}^n (y_i - f(x_i))^2 \quad (27.251)$$

Minimizing the sum of the difference between the observed y_i and the estimates $f(x_i)$ also minimizes their mean, so with this we have also shown that the estimator that minimizes mean square error MSE (see Definition 27.55) is the best estimator according to the likelihood of the produced outcomes.

$$f^* \in F : \frac{1}{n} \sum_{i=1}^n (y_i - f^*(x_i))^2 = \min_{\forall f \in F} \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 \quad (27.252)$$

$$f^* \in F : \text{MSE}(f^*) = \min_{\forall f \in F} \text{MSE}(f) \quad (27.253)$$

The term $(y_i - f(x_i))^2$ is often justified by the statement that large deviations of f from the y -values are punished harder than smaller ones. The correct reason why we minimize the square error, however, is that we maximize the likelihood of the resulting estimator.

At this point, one should also notice that the x_i also could be replaced with vectors $\mathbf{x}_i \in \mathbb{R}^m$ without any further implications or modifications of the equations.

In most practical cases, the set F of possible functions is closely defined. It usually contains only one type of parameterized function, so we only have to determine the unknown parameters in order to find f^* . Let us consider a set of linear functions as example. If we

⁶⁴ http://en.wikipedia.org/wiki/Maximum_likelihood [accessed 2007-07-03]

want to find estimators of the form $F = \{\forall f(x) = ax + b : a, b \in \mathbb{R}\}$, we will minimize Equation 27.254 by determining the best possible values for a and b .

$$\text{MSE}(f(x)|a, b) = \frac{1}{n} \sum_{i=1}^n (ax_i + b - y_i)^2 \quad (27.254)$$

If we now could find a perfect estimator f_p^* and our data would be free of any measurement error, all parts of the sum would become zero. For $n > 2$, this perfect estimator would be the solution of the over-determined system of linear equations illustrated in Equation 27.255.

$$\begin{aligned} 0 &= ax_1 + b - y_1 \\ 0 &= ax_2 + b - y_2 \\ \dots &\dots \\ 0 &= ax_n + b - y_n \end{aligned} \quad (27.255)$$

Since it is normally not possible to obtain a perfect estimator because there are measurement errors or other uncertainties like unknown dependencies, the system in Equation 27.255 often cannot be solved but only minimized.

27.6.2 Best Linear Unbiased Estimators

The Gauss-Markov Theorem⁶⁵ defines BLUEs (best linear unbiased estimators) according to the facts just discussed:

Definition 27.59 (BLUE). In a linear model in which the measurement errors ε_i are uncorrelated and are all normally distributed with an expected value of zero and the same variance, the best linear unbiased estimators (BLUE) of the (unknown) coefficients are the least-square estimators [1740].

Hence, for the best linear unbiased estimator also the same three assumptions (Equation 27.238, Equation 27.239, and Equation 27.240 on page 451) as for the maximum likelihood estimator hold.

27.6.3 Confidence Intervals

Definition 27.60 (Confidence Interval). Unlike point estimators, which approximate a parameter of a data sample with a single value, confidence intervals⁶⁶ (CIs) are estimations that give certain upper and lower boundaries in which the true value of the parameter will be located with a certain, predefined probability. [1741, 1742]

The advantage of confidence intervals is that we can directly derive the significance of the data samples from them – the larger the intervals are, the less reliable is the sample. The narrower confidence intervals get for high predefined probabilities, the more profound, i. e., significant, will the conclusions drawn from them be.

Example

Imagine we run a farm and own 25 chickens. Each chicken lays one egg a day. We collect all the eggs in the morning and weigh them in order to find the average weight of the eggs produced by our farm. Assume our sample contains the values (in g):

⁶⁵ http://en.wikipedia.org/wiki/Gauss-Markov_theorem [accessed 2007-07-03], <http://www.answers.com/topic/gauss-markov-theorem> [accessed 2007-07-03]

⁶⁶ http://en.wikipedia.org/wiki/Confidence_interval [accessed 2007-10-01]

$$A = \{ 120, 121, 119, 116, 115, 122, 121, 123, 122, 120, 119, 122, 121, 120, 119, 121, 123, 117, 118, 121 \} \quad (27.256)$$

$$n = \text{len}(A) = 20 \quad (27.257)$$

From these measurements, we can determine the arithmetic mean $\bar{a}$ and the sample variance s^2 according to Equation 27.60 on page 424 and Equation 27.67 on page 425:

$$\bar{a} = \frac{1}{n} \sum_{i=0}^{n-1} A[i] = \frac{2400}{20} = 120 \quad (27.258)$$

$$s^2 = \frac{1}{n-1} \sum_{i=0}^{n-1} (A[i] - \bar{a})^2 = \frac{92}{19} \quad (27.259)$$

The question that arises now is if the mean of 120 is significant, i. e., whether it likely approximates the expected value of the egg weight, or if the data sample was too small to be representative. Furthermore, we would like to know in which interval the expected value of the egg weights will likely be located. Now confidence intervals come into play. First, we need to find out what the underlying distribution of the random variable producing A as sample output is. In case of chicken eggs, we safely can assume that it is the normal distribution discussed in Section 27.4.2 on page 437. Knowing that, we can now calculate an interval which includes the unknown parameter μ (i. e., the real expected value) with a confidence probability of γ . $\gamma = 1 - a$ is the so-called confidence coefficient and a is the probability that the real value of the estimated parameter lies not inside the confidence interval.

Let us compute the interval including the expected value μ of the chicken egg weights with a probability of $\gamma = 1 - a = 95\%$. Thus, $a = 0.05$. Therefore, we have to pick the right formula from Section 27.6.3 on the following page (here it is Equation 27.272 on the next page) and substitute in the proper values:

$$\mu_\gamma \in \left[\bar{a} \pm t_{1-\frac{a}{2}, n-1} \frac{s}{\sqrt{n}} \right] \quad (27.260)$$

$$\mu_{95\%} \in \left[120 \pm t_{0.975, 19} * \frac{\sqrt{\frac{92}{19}}}{\sqrt{19}} \right] \quad (27.261)$$

$$\mu_{95\%} \in [120 \pm 2.093 * 0.5048] \quad (27.262)$$

$$\mu_{95\%} \in [118.94, 121.06] \quad (27.263)$$

The value of $t_{19, 0.025}$ can easily be obtained from Table 27.12 on page 447 which contains the respective quantiles of Student's t-distribution discussed in Section 27.4.5 on page 445. Let us repeat the procedure in order to find the interval that will contain μ with probabilities $1 - \gamma = 99\% \Rightarrow a = 0.01$ and $1 - \gamma = 90\% \Rightarrow a = 0.1$

$$\mu_{99\%} \in [120 \pm t_{0.995, 19} * 0.5048] \quad (27.264)$$

$$\mu_{99\%} \in [120 \pm 2.861 * 0.5048] \quad (27.265)$$

$$\mu_{99\%} \in [118.56, 121.44] \quad (27.266)$$

$$\mu_{90\%} \in [120 \pm t_{0.95, 19} * 0.5048] \quad (27.267)$$

$$\mu_{90\%} \in [120 \pm 1.729 * 0.5048] \quad (27.268)$$

$$\mu_{90\%} \in [119.13, 120.87] \quad (27.269)$$

$$(27.270)$$

As you can see, the higher the confidence probabilities we specify the larger become the intervals in which the parameter is contained. We can be to 99% sure that the expected value of laid eggs is somewhere between 118.56 and 121.44. If we narrow the interval down to $[119.13, 120.87]$, we can only be 90% confident that the real expected value falls in it based on the data samples which we have gathered.

Some Hand-Picked Confidence Intervals

The following confidence intervals are two-sided, i. e., we determine a range $\tilde{\theta}_\gamma \in [\tilde{\theta}' - x, \tilde{\theta}' + x]$ that contains the parameter θ with probability γ based on the estimate $\tilde{\theta}$. If you need a one-sided confidence interval like $\tilde{\theta}_\gamma \in (-\infty, \tilde{\theta} + x]$ or $\tilde{\theta}_\gamma \in [\tilde{\theta}' - x, \infty)$, you just need to replace $1 - \frac{a}{2}$ with $1 - a$ in the equations.

Expected Value of a Normal Distribution $N(\mu, \sigma^2)$

With knowing the variance σ^2 : If the exact variance σ^2 of the distribution underlying our data samples is known, and we have an estimate of the expected value μ by the arithmetic mean $\bar{a}$ according to Equation 27.60 on page 424, the two-sided confidence interval (of probability γ) for the expected value of the normal distribution is:

$$\mu_\gamma \in \left[\bar{a} \pm z \left(1 - \frac{a}{2} \right) \frac{\sigma}{\sqrt{n}} \right] \quad (27.271)$$

Where $z(y) \equiv \text{probit}(y) \equiv \Phi^{-1}y$ is the y -quantile of the standard normal distribution (see Definition 27.49 on page 439) which can for example be looked up in Table 27.7.

With estimated sample variance s^2 : Often, the true variance σ^2 of an underlying distribution is not known and instead estimated with the sample variance s^2 according to Equation 27.67 on page 425. The two-sided confidence interval (of probability γ) for the expected value can then be computed using the arithmetic mean $\bar{a}$ and the estimate of the standard deviation $s = \sqrt{s^2}$ of the sample and the $t_{n-1, 1-\frac{a}{2}}$ quantile of Student's t-distribution which can be looked up in Table 27.12 on page 447.

$$\mu_\gamma \in \left[\bar{a} \pm t_{1-\frac{a}{2}, n-1} \frac{s}{\sqrt{n}} \right] \quad (27.272)$$

Variance of a Normal Distribution

The two-sided confidence interval (of probability γ) for the variance of a normal distribution can be computed using sample variance s^2 and the $\chi^2(p, k)$ -quantile of the χ^2 distribution which can be looked up in Table 27.10 on page 444.

$$\sigma_\gamma^2 \in \left[\frac{(n-1)s^2}{\chi^2(1-\frac{a}{2}, n-1)}, \frac{(n-1)s^2}{\chi^2(\frac{a}{2}, n-1)} \right] \quad (27.273)$$

Success Probability p of a $B(1, p)$ Binomial Distribution

The two-sided confidence interval (of probability γ) of the success probability p of a $B(1, p)$ binomial distribution can be computed as follows:

$$p_\gamma \in \left[\frac{n}{n + z_{1-\frac{a}{2}}^2} \left(\bar{a} + \frac{1}{2n} z_{1-\frac{a}{2}}^2 \pm z_{1-\frac{a}{2}} \sqrt{\frac{\bar{a}(1-\bar{a})}{n} + \left(\frac{1}{2n} z_{1-\frac{a}{2}}^2 \right)^2} \right) \right] \quad (27.274)$$

Expected Value of an Unknown Distribution with Sample Variance

The two-sided confidence interval (of probability γ) of the expected value EX of an unknown distribution with an unknown real variance D^2X can be determined using the arithmetic mean $\bar{a}$ and the sample variance s^2 if the sample data set contains more than $n = 50$ elements.

$$EX_\gamma \in \left[\bar{a} \pm z \left(1 - \frac{a}{2} \right) \frac{s}{\sqrt{n}} \right] \quad (27.275)$$

27.6.4 Density Estimation

In this section we discuss density estimation⁶⁷ techniques [1743, 1744]. Density estimation is often used by global optimization algorithms in order to test whether a certain region of the search space has already been explored sufficiently and where to concentrate the further search efforts.

Definition 27.61 (Density Estimation). A density estimation $\rho(a)$ approximates an unobservable probability density function $f_X(a)$ (see Section 27.1.8 on page 422) using a set of sample data A .

$$\rho(a) \approx f_X(a) \quad (27.276)$$

$$\rho : A \rightarrow \mathbb{R}^+ \quad (27.277)$$

Histograms

TODO

The k^{th} Nearest Neighbor Method

Definition 27.62 (k^{th} Nearest Neighbor Distance). The k^{th} nearest neighbor distance function $\text{dist}_{nn,k}^\rho$ denotes the distance of one element a to its k^{th} nearest neighbor in the set of all elements A . It relies on a distance measure (here called dist) to compute the element distances. See Section 28.1 on page 469 for more details on distance measures.

$$\text{dist}_{nn,k}^\rho(\text{dist}, a, A) = \text{dist}(a, a_k) : |\forall b \in A : \text{dist}(a, b) < \text{dist}(a, a_k)| = k - 1 \quad (27.278)$$

Using the k^{th} nearest neighbor method [1745], the probability density function of an element a is estimated by its distance to its k^{th} nearest neighbor a_k in the test set A (with $k < |A|$). Most often, the k^{th} nearest neighbor distance measure internally uses the Euclidian distance measure $\text{dist}_{eucl} \equiv \text{dist}_{n,2}$ (see Definition 28.8 on page 470), but theoretically any other one of the distance measures presented in Section 28.1 could also be applied. Normally, k is chosen to be $\sqrt{|A|}$

$$\rho_{nn,k}(a, A) = \frac{k}{2 |A| \text{dist}_{nn,k}^\rho(\text{dist}, a, A)} \quad (27.279)$$

Crowding Distance

Crowding distance [1746] treats every element $a \in A$ as n -dimensional vector (where each dimension will represent an objective subject to optimization in the context of this book). The crowding distance is not a distance measure as its name may suggest, but a base for a density estimate. When computing the crowding distance of an element a we consider every single dimension i of the element a separately. For each of its dimensions, we determine the nearest neighbor to the left a^l and the nearest neighbor to the right a^r . The crowding distance of the element a in the dimension i is then $a_i^r - a_i^l$, the distance of the (objective) values of the right and left neighbors of a in the dimension i . This distance is normalized so that the maximum crowding distance of all elements in A in any dimension is 1. If an element has no left or no right neighbor in this dimension, meaning that it is situated on either end of the spectrum represented by all elements in the sample A , its crowding distance in the dimension is also set to 1.

⁶⁷ http://en.wikipedia.org/wiki/Density_estimation [accessed 2007-07-03]

The original source [1746] does not mention normalization explicitly and sets the crowding distance of edge elements to ∞ , which is both problematic. If no normalization is performed, dimensions with large crowding distances will outweigh those with smaller values – they will play no role in the crowding density value finally computed. With normalization, each dimension has the same weight. If the crowding distance of edge elements is set to ∞ , they will have a very outstanding position in A which could influence processes relying on the crowding distance in a very strong way.

The total crowding distance of an element a is the sum of its distance values corresponding to each dimension. Algorithm 27.1 on the facing page computes a function $\text{dist}_{cr}^\rho(a, A)$ which relates each element a of the set A to its crowding distance. In this algorithm, we consider dist_{cr}^ρ to be some sort of lookup-table instead of a mathematical function. Therefore, we can build it iteratively by summing up the distance values dimension-wise. Since computing the crowding distance can be performed best by sorting the individuals according to their values in the single dimensions, we define the comparator function⁶⁸ cmp_i as follows:

$$\text{cmp}_{cr,i}(a, b) = \begin{cases} -1 & \text{if } a_i < b_i \\ 0 & \text{if } a_i = b_i \\ 1 & \text{otherwise} \end{cases} \quad \forall a, b \in A, \forall i \in [0, |a|] \quad (27.280)$$

The crowding distance can now be used as density estimate whereas individuals with large crowding distance values are in a sparsely covered region while small values of dist_{cr}^ρ indicate dense portions of A . A density estimate derived from the crowding distance will therefore be inversely proportional to it. Hence, we define the density measure ρ_{cr} as the difference of 1 and $\text{dist}_{cr}^\rho(a, A)$ divided by the vector dimensions $n = |a|$, obtaining a value in $[0, 1]$ that is big if a is in crowded region and small if it is situated in a sparsely covered area of A . It should be noted that this density estimate is mathematically not fully sound since it only displays the crowding information.

$$\rho_{cr}(a, A) = 1 - \frac{\text{dist}_{cr}^\rho(a, A)}{n} \quad (27.281)$$

Parzen Window / Kernel Density Estimation

Another density estimation is the Parzen [1747] window method⁶⁹, also called kernel density estimation.

⁶⁸ Comparator functions were introduced in Definition 1.15 on page 19.

⁶⁹ http://en.wikipedia.org/wiki/Parzen_window [accessed 2007-07-03]

Algorithm 27.1: $\text{dist}_{cr}^{\rho}(\dots, A) \leftarrow \text{computeCrowdingDistance}(a, A)$

Input: A : the set of sample data
Data: dd : a list used as store for the crowding distances of the single dimensions
Data: A_s : the list representation of A
Data: dim : the dimension counter
Data: j : the element counter
Data: max : the maximum crowding distance of the current dimension
Output: $\text{dist}_{cr}^{\rho}(\dots, A)$: the crowding distance function

```

1 begin
2    $dd \leftarrow \text{createList}(\text{len}(A), 0)$ 
3    $dd[0] \leftarrow 1$ 
4    $dd[\text{len}(A)-1] \leftarrow 1$ 
5    $A_s \leftarrow \text{setToList}(A)$ 
6    $dim \leftarrow n$ 
7   while  $dim > 0$  do
8      $A_s \leftarrow \text{sortList}_a(A, \text{cmp}_{dim})$ 
9      $max \leftarrow 0$ 
10     $j \leftarrow \text{len}(A) - 2$ 
11    while  $j > 0$  do
12       $dd[j] \leftarrow A_s[j+1]_{dim} - A_s[j-1]_{dim}$ 
13      if  $dd[j] > max$  then  $max \leftarrow dd[j]$ 
14       $j \leftarrow j - 1$ 
15    if  $max > 0$  then
16       $j \leftarrow \text{len}(A) - 2$ 
17      while  $j > 0$  do
18         $dd[j] \leftarrow \frac{dd[j]}{max}$ 
19         $j \leftarrow j - 1$ 
20     $j \leftarrow \text{len}(A) - 1$ 
21    while  $j \geq 0$  do
22       $\text{dist}_{cr}^{\rho}(A_s[j], A) \leftarrow \text{dist}_{cr}^{\rho}(A_s[j], A) + dd[j]$ 
23       $j \leftarrow j - 1$ 
24     $dim \leftarrow dim - 1$ 
25  return  $\text{dist}_{cr}^{\rho}(\dots, A)$ 
26 end
  
```

TODO

27.7 Generating Random Numbers

Definition 27.63 (Random Number). Random numbers are the values taken on by a random variable. A random number generator⁷⁰ produces a sequence $r = (r_1, r_2, \dots)$ of random numbers r_i as result of independent repetitions of the same random experiment.

Since the numbers r_i are all produced by the same random experiment, they approximate a certain random distribution according to the law of large numbers (see Section 27.2.8 on page 429).

For true random number generators, there exists no function or algorithm $f(i) = r_i$ or $f(r_{i-n+1}, r_{i-n+2}, \dots, r_i) = r_{i+1}$ that can produce this sequence in a deterministic manner with or without knowledge of the random numbers previously returned from the generator. Such

⁷⁰ http://en.wikipedia.org/wiki/Random_number_generator [accessed 2007-07-03], http://en.wikipedia.org/wiki/Random_number_generation [accessed 2007-07-03]

behavior can be achieved by obtaining the numbers r_i from measurements of a physical process, for instance. Today, there exist many such so-called hardware random number generators⁷¹ [1748, 1749, 1750, 1751].

Of course, most computers are not equipped with special hardware for random number production, although some standard devices can be utilized for that purpose. One could, for example, measure the white noise of soundcards or the delays between the user's keystrokes. Such methods have the drawback that they require the presence of and access to such components. Furthermore, the speed of them is limited since you cannot produce random numbers faster than the recording speed of the soundcard or faster than the user is typing.

27.7.1 Generating Pseudorandom Numbers

In security-sensitive areas like cryptography, we need true random numbers [1752, 1751, 1753]. For normal PC applications and most scientific purposes, pseudorandom number generators⁷² are sufficient.

The principle of pseudorandom number generators is to produce a sequence of numbers $r = (r_1, r_2, \dots, r_i), r_j \in R \forall j \in \mathbb{N}, R \subseteq \mathbb{R}$ which are not *obviously* interdependent, i. e., if knowing a number r_i there is not simple way to find out the value of r_{i+1} .

Of course, since the values r_i are no real random numbers, there is an algorithm or function $f : V \rightarrow R \times V$ where R is the set of possible numbers and V is the space of some internal variables. These internal variables are referred to as seed and normally change whenever a new number is produced. Often, the seed is initialized with either a true random number or the current system time. In the first case, it is also practicable to re-initialize the seed from time to time with new true random values.

Pseudorandom numbers are attractive to all not security-critical applications where we need some sort of unpredictable behavior. They are often used in games or simulations, since they usually can be generated much quicker than true random numbers. On the other hand, especially in scientific applications the "degree" of randomness is very important. There are many incidents, for example in physical simulation, where the inappropriate use of pseudorandom number generators of poor quality lead to wrong conclusions [1754, 1755, 1756]. It should be noted that there also exist cryptographically secure pseudorandom number generators⁷³ which create pseudo-random number that are of especially high quality.

There exists a variety of algorithms that generate pseudorandom numbers [1757, 1758, 1759, 1760] and many implementations for different programming languages and architectures [1761, 1762, 1763]. It is even possible to evolve pseudorandom number generators using Genetic Programming, as shown, for instance, by Koza [1764].

Linear Congruential Generator (LCG)

The linear congruential generator⁷⁴ (LCG) was first proposed by Lehmer [1765] and is one of the most frequently used and simplest pseudo random number generators [1766]. It updates an internal integer number $v \in V = (0 \dots (m - 1)), m \in \mathbb{N}$ in each step according to Equation 27.282. The modulus m is a natural number which defines the maximum number of values v can take on. a and b are both constants. Therefore, v will periodically take on the same values – at most after m steps. The pseudorandom numbers r_i are approximately uniformly distributed in the interval $[0, m)$ (see Section 27.3.1) and can be computed as proposed in Equation 27.283.

⁷¹ http://en.wikipedia.org/wiki/Hardware_random_number_generator [accessed 2007-07-03]

⁷² http://en.wikipedia.org/wiki/Pseudorandom_number_generator [accessed 2007-07-03], <http://en.wikipedia.org/wiki/Pseudorandomness> [accessed 2007-07-03]

⁷³ http://en.wikipedia.org/wiki/Cryptographically_secure_pseudorandom_number_generator [accessed 2007-07-03]

⁷⁴ http://en.wikipedia.org/wiki/Linear_congruential_generator [accessed 2007-07-03]

$$v_i = (av_{i-1} + b) \bmod m \quad (27.282)$$

$$r_i = \frac{v_i}{m} \quad (27.283)$$

If the full period can really be reached depends a lot on the values of the parameters a , b , and m . There are many constellations known where only a small fraction of the period m is utilized [1767]. In order to produce the full period, the following requirements should be met according to Wikipedia [2].

- b and m are relatively prime
- $a - 1$ is divisible by all prime factors of m
- $a - 1$ is a multiple of 4 if m is a multiple of 4
- $m > \max\{a, b, v_0\}$
- $a > 0, b > 0$

Good standard values for the constants are $a = 1\,664\,525$, $b = 1\,013\,904\,223$, and $m = 2^{32}$. One of the widest spread realizations of LCGs has been outlined by Knuth [1768]. In Java, the class `java.util.Random` uses this approach with the settings $a = 25\,214\,903\,917$, $b = 11$, and $m = 2^{48}$.

27.7.2 Random Functions

Definition 27.64 (Random Function). In the context of this book, we define a random function `random` as a construct that eases the utilization of random numbers and random variables. It represents access to a random process, an infinite sequence of random variables X_i all distributed according to the same distribution function. Starting with X_1 , each time a random function is evaluated, it returns the value of the next random variable in the sequence $i = 1, 2, 3, \dots$

Definition 27.65 (Uniformly Distributed Random Number Generator). We define the function `randomu( $\check{r}, \hat{r}$ )` to draw uniformly distributed (see Section 27.4.1 on page 436) random numbers from the interval with the boundaries $\check{r}$ (inclusively) and $\hat{r}$ (exclusively). The parameter-less function `randomu()` will return a uniformly distributed number from the interval spanning from 0 inclusively to 1 exclusively.

$$\text{random}_u(\check{r}, \hat{r}) \in [\check{r}, \hat{r}) \subseteq \mathbb{R}, \check{r}, \hat{r} \in \mathbb{R}, \check{r} < \hat{r} \quad (27.284)$$

$$\text{random}_u() \equiv \text{random}_u(0, 1) \quad (27.285)$$

The `randomu()`-function can be realized with the linear congruential pseudorandom number generators that we have just discussed in Section 27.7.1, for example.

Definition 27.66 (Normally Distributed Random Number Generator). We define the function `randomn( $\mu, \sigma^2$ )` to generate normally distributed (see Section 27.4.2 on page 437) random numbers with the expected value μ and the variance σ^2 . The parameter-less function `randomn()` will return a standard normally distributed number (with $\mu = 0$ and $\sigma^2 = 1$).

$$\text{random}_n(\mu, \sigma^2) \sim N(\mu, \sigma^2) \quad (27.286)$$

$$\text{random}_n() \equiv \text{random}_n(0, 1) \quad (27.287)$$

Cut-off Random Functions

We often use random processes and random functions to model or simulate a certain features of a real system. If we, for example, simulate a chicken farm, we might be interested in the size of the eggs laid by the hens. We can assume this weight to be normally distributed⁷⁵ around some mean μ with a variance $\sigma^2 \neq 0$. In the simulation, a series of egg weights is created simply by drawing subsequent such random numbers by calling $\text{random}_n(\mu, \sigma^2)$ repeatedly. Although the normal distribution is a good model for egg weights, it has a serious drawback: no matter how we chose μ or σ , there is still a positive probability of drawing zero, negative, or extremely large ($> 10t$) weights. In reality however, such things could have not yet been documented.

What we need here is a cut-off mechanism for our random function $\text{random}_n(\mu, \sigma^2)$ that still preserves as many of its properties as possible. Given any random function random the function $\text{random}_l(\text{random}, \text{low}, \text{high})$, defined as Algorithm 27.2, ensures that $\text{low} \leq \text{random}_l(\text{random}, \text{low}, \text{high}) < \text{high}$.

Algorithm 27.2: $r \leftarrow \text{random}_l(\text{random}, \text{low}, \text{high})$

Input: random : a random function (maybe with further implicit parameters)
Input: $\text{low} \in \mathbb{R}$: the inclusive, lower bound of the random result
Input: $\text{high} \in \mathbb{R}, \text{low} < \text{high}$: the exclusive, upper bound of the random result
Data: r : the intermediate random value
Output: r : a value returned by random with $\text{low} \leq r < \text{high}$

```

1 begin
2   repeat
3     |  $r \leftarrow \text{random}()$ 
4   until  $(r \geq \text{low}) \wedge (r < \text{high})$ 
5   return  $r$ 
6 end
```

27.7.3 Converting Random Numbers to other Distributions

There are occasions where random numbers of a different distribution than available are needed. We could, for example, have a linear congruential generator for uniformly distributed random numbers like elaborated in Section 27.7.1 but may need normally distributed values.

Uniform Distribution $\rightarrow$ Uniform Distribution

If we have random numbers r_i distributed uniformly in the interval $[a_1, b_1)$ and need random numbers s_i uniformly distributed in the interval $[a_2, b_2)$, they can be converted really simple according to

$$s_i = a_2 + (b_2 - a_2) \frac{r_i - a_1}{b_1 - a_1} \quad (27.288)$$

Uniform Distribution $\rightarrow$ Normal Distribution

In order to transform random numbers which are uniformly distributed in the interval $[0, 1)$ to standard-normally distributed random numbers ($\mu = 0, \sigma^2 = 1$), we can apply the Box-Muller⁷⁶ transformation [1769]. This approach creates two standard-normally distributed

⁷⁵ see Section 27.4.2 on page 437

⁷⁶ http://en.wikipedia.org/wiki/Box_muller [accessed 2007-07-03]

random numbers n_1, n_2 from two random numbers r_1, r_2 which are uniformly distributed in $[0, 1)$ at once according to Equation 27.289. In both formulas, the terms $\sqrt{-2 \ln r_1}$ and $2\pi r_2$ are used. The performance can be increased if both terms are computed only once and reused.

$$\begin{aligned} n_1 &= \sqrt{-2 \ln r_1} \cos(2\pi r_2) \\ n_2 &= \sqrt{-2 \ln r_1} \sin(2\pi r_2) \end{aligned} \quad (27.289)$$

The polar form of this method, illustrated as Algorithm 27.3, is not only faster, but also numerically more robust [1730]. It creates two independent random numbers uniformly distributed in $[-1, 1)$ and computes their product w . This is repeated until $w \in (0, 1)$. With this value, we now can compute two independent, standard-normally distributed random numbers. Effectively, we have traded a trigonometric operation and a multiplication against a division compared to the original method in Equation 27.289. The implementation of this algorithm is discussed in [1768] which is the foundation of the method `nextGaussian` of the Java-class `java.util.Random`.

Algorithm 27.3: $(n_1, n_2) \leftarrow \text{random}_{n,p^2}()$

Data: n_1, n_2 : the intermediate and result variables

Data: w : the polar radius

Output: (n_1, n_2) : a tuple of two independent values $n_1 \sim N(0, 1), n_2 \sim N(0, 1)$

```

1 begin
2   repeat
3      $n_1 \leftarrow \text{random}_u(-1, 1)$ 
4      $n_2 \leftarrow \text{random}_u(-1, 1)$ 
5      $w \leftarrow (n_1 * n_1) + (n_2 * n_2)$ 
6   until  $(w > 0) \wedge (w < 1)$ 
7    $w \leftarrow \sqrt{\frac{-2 \ln w}{w}}$ 
8   return  $(n_1 * w, n_2 * w)$ 
9 end
```

Normal Distribution $\rightarrow$ Normal Distribution

With Equation 27.290, a normally distributed random number $n_1 \sim N(\mu_1, \sigma_1^2)$ can be transformed to another normally distributed random number $n_2 \sim N(\mu_2, \sigma_2^2)$.

$$n_2 = \mu_2 + \sigma_2 * \frac{n_1 - \mu_1}{\sigma_1} \quad (27.290)$$

Uniform Distribution $\rightarrow$ Exponential Distribution

With Equation 27.291, a random number r uniformly distributed in the interval $(0, 1)$ (0 is excluded) can be transformed into an exponentially distributed random number $s \sim \text{exp}(\lambda)$.

$$s = \frac{-\ln r}{\lambda} \quad (27.291)$$

Exponential Distribution → Exponential Distribution

With Equation 27.292, an exponentially distributed random number $r_1 \sim \exp(\lambda_1)$ can be transformed to an exponentially distributed number $r_2 \sim \exp(\lambda_2)$.

$$r_2 = \frac{\lambda_1}{\lambda_2} r_1 \quad (27.292)$$

Uniform Distribution → Bell-shaped Distribution

The bases of many numerical optimization algorithms is the modification of a value x by adding some random number to it. If the probability density function of the underlying distribution producing number is symmetrically bell-shaped, the result of the additive modification will be smaller or larger than x with the same probability. Results which are close to x will be more likely than such that are very distant. One example for such a distribution is the normal distribution. Another example is the bell-shaped random number generator used by Worakul et al. [1293; 1294], defined here as Algorithm 27.4. It is algorithmically close to the polar form of the Box-Muller transform for the normal distribution (see Algorithm 27.3 on the preceding page) but differs in the way the internal variable w is created. The function $\text{random}_{bs}(\mu, \sigma)$ creates a new random number according to this distribution, with an expected value μ and the standard deviation σ .

Algorithm 27.4: $y \leftarrow \text{random}_{bs}(\mu, \sigma)$

Input: μ : the mean value of the bell-shaped distribution

Input: σ : the approximate standard deviation of the bell-shaped distribution

Data: w : a uniformly distributed random number $w \in (0, 1)$

Output: y : a bell-shaped distributed random number

```

1 begin
2   repeat
3      $w \leftarrow \text{random}_u()$ 
4   until  $(w > 0) \wedge (w < 1)$ 
5    $y \leftarrow \mu + \sigma * 0.5513 * \ln\left(\frac{1-w}{w}\right)$ 
6   return  $y$ 
7 end
```

You may have wondered about the factor 0.5513 in the algorithm. This number “normalizes” the standard deviation of the bell-shaped distribution, since $D^2\left[r(y) = \ln\left(\frac{1-y}{y}\right)\right] \neq 1$. We can show this by first determining the cumulative distribution function $F_X(x)$ for $r(y)$ in Equation 27.295 and then differentiating in order to obtain the probability density function $f_X(x)$ in Equation 27.297.

```

1 long i, max;
2 double sum2, v;
3
4 max    = 10000000;
5 sum2   = 0;
6 v      = 0;
7
8 // distribution is symmetric -> iterate one wing
9 for (i = (max>>1); i < max; i++) {
10     v = Math.log(((double) (max - i)) / ((double) i));
11     sum2 += (v * v); //sum up the squares of the single terms
12 }
13
14 System.out.print(sum2 / (max - (max>>1)));

```

Listing 27.1: Approximating D^2X of $r(y)$.

$$F_X(x) \equiv r^{-1}(0, 1) \quad (27.293)$$

$$x = r(y) = \ln \left(\frac{y}{1-y} \right) \quad (27.294)$$

$$F_X(x) = y = \frac{e^x}{1+e^x} \quad (27.295)$$

$$f_X(x) = F_X(x) \frac{dx}{dy} \quad (27.296)$$

$$\left(\frac{e^x}{1+e^x} \right) \frac{dx}{dy} = \frac{e^x (1+e^x) - e^x (e^x)}{(1+e^x)^2}$$

$$f_X(x) = \frac{e^x}{(1+e^x)^2} \quad (27.297)$$

Unfortunately, here it stops. We can neither apply Equation 27.56 on page 424 or Equation 27.63 on page 425 in order to determine the expected value or the variance, since both will result in integrals that the author⁷⁷ cannot compute. However, it is easy to see that $EX = 0$, since $r(y)$ is point symmetric around 0.5. The value $D^2X \approx 3.28984$ I can only determine numerically with the small Java program Listing 27.1 which bases on the idea that we can assume the uniform random numbers to be uniformly distributed in $(0, 1)$ (of course). Hence we can simulate a “complete sample” by iterating over $i = 1$ to $T-1$ and take i/T as input for $r(y)$. Since we step over all i from 1 to $T-1$, this resembles a uniform distribution and also leaves away the special cases $y = 0$ ($\sim i=0$) and $y = 1$ ($\sim i=T$). Furthermore, we can skip half of the steps since our distribution is symmetric. Well, $EX = 0$ if $\mu = 0$ and therefore we can simplify $D^2X = E[X] - (EX)^2$ (see Equation 27.61 on page 424) to $D^2X = E[X]$.

This method is, of course, very crude and subject to numerical errors in the floating point computations. However, with $D^2X \approx 3.28984$ and $DX = \sqrt{D^2X} \approx 1.8138$ we know that we have to scale $r(y)$ by $\frac{1}{DX} \approx 0.5513$ (see Equation 27.65 on page 425) so the standard deviation the bell-shaped distribution $\text{random}_{bs}(\mu, \sigma)$ will become $D[\text{random}_{bs}(\mu, \sigma)] \approx \sigma$.

⁷⁷ Yes. I suck in maths.

27.8 List of Functions

27.8.1 Gamma Function

Definition 27.67 (Gamma Function). The Gamma function⁷⁸ $\Gamma : \mathbb{C} \mapsto \mathbb{R}$ is the extension of the factorial (see Definition 27.8 on page 418) to the real and complex numbers. For complex numbers $z \in \mathbb{C}$ with a positive real part $\operatorname{Re}(z) > 0$ it is defined as:

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt \quad (27.298)$$

Furthermore, the following equations hold for the gamma function.

$$\Gamma(z+1) = z\Gamma(z) \quad (27.299)$$

$$\Gamma(1) = 1 \quad (27.300)$$

$$\Gamma(z) = (z-1)! \quad \forall z \in \mathbb{N} \quad (27.301)$$

$$\Gamma(z) = \lim_{n \rightarrow \infty} \frac{n! n^z}{z(z+1)\dots(z+n)} \quad (27.302)$$

$$\Gamma(z) = \frac{e^{\gamma z}}{z} \prod_{n=1}^{\infty} \left(1 + \frac{z}{n}\right)^{-1} e^{\frac{z}{n}} \quad (27.303)$$

γ in Equation 27.303 denotes the Euler-Mascheroni constant⁷⁹.

$$\gamma = \lim_{n \rightarrow \infty} \left[\left(\sum_{k=1}^n \frac{1}{k} \right) - \log n \right] = \int_0^\infty \left(\frac{1}{[x]} - \frac{1}{x} \right) dx \quad (27.304)$$

$$\approx 0.57721566490153286060651209008240243 \dots \quad (27.305)$$

⁷⁸ http://en.wikipedia.org/wiki/Gamma_function [accessed 2007-09-30]

⁷⁹ http://en.wikipedia.org/wiki/Euler-Mascheroni_constant [accessed 2007-09-30]

Clustering

Clustering algorithms¹ divide a dataset into several disjoint subsets. All elements in such a subset share common features like, for example, spatial proximity. Clustering has many different applications like:

- Data Mining and Data Analysis [1770, 1771, 1772, 1773],
- Information Processing and Information Management [1774, 1775, 1776, 1777],
- Pattern Recognition [1778, 1779, 1780],
- Image Processing [1566, 1781], and
- Medicine [1782, 1783, 1784].

Definition 28.1 (Clustering). Clustering is the unsupervised classification of patterns (observations, data items, or feature vectors) into groups (clusters) [1785]. With clustering, one dataset is partitioned into subsets (clusters), so that the data in each subset (ideally) share some common trait - often proximity according to some defined distance measure. Figure 28.1 illustrates a possible result C of the application of a clustering algorithm to a set A of elements with two features.




Fig. 28.1: A clustering algorithm applied to a two-dimensional dataset A .

In the field of global optimization there is another application for clustering algorithms. For many problems the set of optimal solutions X^* is very large or even infinite. An optimization algorithm then cannot be able to store or return it on the whole. Therefore, clustering techniques are often used in order to reduce the optimal set while not losing its characteristics – the diversity of the individuals included in the “current optimal set” is preserved, just their number is reduced. Especially in elitist evolutionary algorithms (see Definition 2.4 on page 75) which maintain an archive of the best individuals currently known.

Data clustering algorithms are either hierarchical or partitional. A hierarchical algorithm uses previously established clusters to successively find new clusters. The result of such an

¹ http://en.wikipedia.org/wiki/Data_clustering [accessed 2007-07-03]

algorithm is a hierarchy of clusters. Partitional algorithms, on the other hand, determine all clusters at once. In the context of this book, we do only need the division of a set into clusters – a hierarchy of this division is unnecessary.

There also exist so-called fuzzy clustering² [1786, 1787] methods that do not create clear divisions but assign a vector of probabilities to each element. This vector contains a component for each cluster which denotes the probability of the element to belong to it. Again, in the context of this book, we only regard clustering algorithms that group each data element to exactly one single cluster. Therefore, we define a clustering algorithm as follows:

Definition 28.2 (Clustering Algorithm). A clustering algorithm $C = \text{cluster}(A)$ constructs a set C consisting of elements which are disjoint subsets of a set A and, if united, cover A completely (see also Figure 28.1).

$$\begin{aligned} C = \text{cluster}(A) \Rightarrow & \forall c \in C, \forall a \in c \Rightarrow a \in A \wedge \\ & \forall c_1 \neq c_2 \wedge c_1, c_2 \in C \Rightarrow c_1 \cap c_2 = \emptyset \wedge \\ & \forall a \in A \exists c \in C : a \in c \end{aligned} \quad (28.1)$$

$$\text{deduced: } \bigcup_{c \in C} c = A \quad (28.2)$$

$$\text{deduced: } C \subset \mathcal{P}(A) \quad (28.3)$$

For the last deduced formula see the definition of the power set $\mathcal{P}$, Definition 26.9 on page 410.

There is however one important fact that must not be left unsaid here: Although we define clustering algorithms in terms of sets for simplicity, they are actually applied to lists. A set can contain the *same* element only once, hence $\{1, 2, 1\} = \{1, 2\}$. A clustering algorithm however may receive an input A that contains multiple *equal* elements. This is our little dirty backdoor here, we consider $A = \{a_1, a_2, \dots, a_n\}$ as the input set and allow its elements to have *equal values*, such as $a_1 = 1$, $a_2 = 2$, and $a_3 = 1$. When performing the clustering, we only consider the symbols $a_1 \dots a_n$. This allows us to use straightforward and elegant set-based definitions as done in Definition 28.2 without loss of generality.

Definition 28.3 (Partitions in Clustering). We define the set $\mathfrak{C}$ of all possible partitions of A into clusters C . Furthermore, the subset $\mathfrak{C}_k \subseteq \mathfrak{C}$ is the set of all partitions of A into k clusters. The number of possible configurations $\mathfrak{C}_k$ for any given k equals the Sterling number $S(|A|, k)$ [1788].

$$\begin{aligned} \forall C \in \mathfrak{C} \Leftrightarrow & \forall c \in C, \forall a \in c \Rightarrow a \in A \wedge \\ & \forall c_1, c_2 \in C \Rightarrow c_1 \cap c_2 = \emptyset \wedge \\ & \forall a \in A \exists c \in C : a \in c \end{aligned} \quad (28.4)$$

$$C \in \mathfrak{C}_k \Leftrightarrow C \in \mathfrak{C} \wedge |C| = k \quad (28.5)$$

$$|\mathfrak{C}_k| = S(|A|, k) = \frac{1}{k!} \sum_{i=1}^k (-1)^{k-i} \binom{k}{i} i^n \quad (28.6)$$

$$|\mathfrak{C}| = \sum_{k=1}^n |\mathfrak{C}_k| = \sum_{k=1}^n S(|A|, k) \quad (28.7)$$

On the elements a of the set A which are subject to clustering, we impose an simple restriction: Although we allow any sort of elements a in A , we assume that to each such element a there is assigned exactly one single $\alpha(a) \in \mathbb{R}^n$. In other words, there exists a

² http://en.wikipedia.org/wiki/Fuzzy_clustering [accessed 2007-07-03]

function $\alpha : A \mapsto \mathbb{R}^n$ which relates the features of each element a of A to a vector of real numbers. This allows us to apply distance metrics and such and such.

In the context of global optimization, the elements a would for example be the solution candidates like evolved programs in the population A and the function $\alpha(a)$ then would correspond to the values of their objective functions $f \in F$.

From now on, we will be able to treat the elements a like vectors of real numbers (if needed) without loss of generality. Note that even though we assume that there exists a binary relation which assigns a real vector to each element of A , this is not necessarily the case for the opposite direction. Picking up the previous example it is most probably not likely to have one program for each possible combination of objective values.

Definition 28.4 (Centroid). The centroid³ [1789] of a cluster is its center, the arithmetic mean of all its points to put it plain and simple.

$$\text{centroid}(c) = \frac{1}{|c|} \sum_{\forall a \in c} a \quad (28.8)$$

28.1 Distance Measures

Each clustering algorithm needs some form of distance measuring, be it between two elements or between two clusters. Therefore we define the prototype of a distance measurement function as follows:

Definition 28.5 (Distance Measure). A distance measurement function dist rates the distance between two elements of the same type (set) as positive real number which is the bigger the bigger the distance between the two elements is.

$$\text{dist}(a, b) \in \mathbb{R}^+, a, b \in A \quad (28.9)$$

28.1.1 Distance Measures for Strings of Equal Length

Definition 28.6 (Hamming Distance). For two tuples a and b of the same length, the Hamming [627] distance⁴ $\text{dist}_{Ham}(a, b)$ is defined as the number of locations in which a and b differ.

$$\text{dist}_{Ham}(a, b) = |\{i : a[i] \neq b[i], \forall 0 \leq i < |a|\}| \quad \forall a, b : \text{len}(a) = \text{len}(b) \quad (28.10)$$

The Hamming distance is used in many error-correction schemes, since it also equals to the number of single substitutes required to change one string into another one. The Hamming distance of 100101 and 101001 is 2 whereas the Hamming distance of `Hello World` and `Hello Earth` is 5.

28.1.2 Distance Measures for Real-Valued Vectors

As already mentioned in Chapter 28, we assume that there is a real-values vector in $\mathbb{R}^n$ assigned to each element $a \in A$ by an implicit $\alpha : A \mapsto \mathbb{R}^n$ -function. Therefore, the distance measures introduced here can be used for all A subject to clustering.

Definition 28.7 (Manhattan Distance). The Manhattan distance⁵ $\text{dist}_{Man}(\mathbf{a}, \mathbf{b})$ denotes the sum of the absolute distances of the coordinates of the two vectors $\mathbf{a}$ and $\mathbf{b}$.

$$\text{dist}_{Man}(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^n |\mathbf{a}[i] - \mathbf{b}[i]| \quad \forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^n \quad (28.11)$$

³ <http://en.wikipedia.org/wiki/Centroid> [accessed 2007-07-03]

⁴ http://en.wikipedia.org/wiki/Hamming_distance [accessed 2007-07-03]

⁵ http://en.wikipedia.org/wiki/Manhattan_distance [accessed 2007-07-03]

Thus, the Manhattan distance of $(1, 2, 3)^T$ and $(3, 2, 1)^T$ is 4.

Definition 28.8 (Euclidian Distance). The Euclidian distance⁶ $\text{dist}_{eucl}(\mathbf{a}, \mathbf{b})$ is the "ordinary" distance of two points (denoted by the two vectors $\mathbf{a}$ and $\mathbf{b}$) in Euclidian space. This value is obtained by applying of the Pythagorean theorem⁷.

$$\text{dist}_{eucl}(\mathbf{a}, \mathbf{b}) = \sqrt{\sum_{i=1}^n (\mathbf{a}[i] - \mathbf{b}[i])^2} \quad \forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^n \quad (28.12)$$

Therefore, the Euclidian distance of $(1, 2, 3)^T$ and $(3, 2, 1)^T$ is $\sqrt{8}$.

Definition 28.9 (Norm). A vector norm⁸, denoted by $\|\mathbf{a}\|$ is a function which assigns a positive length or size to all vectors $\mathbf{a}$ in a vector space (or set) $A \subseteq \mathbb{R}^n$, other than the zero vector.

Some common norms of the element $\mathbf{a} \in A \subseteq \mathbb{R}^n$ are:

- The Manhattan norm⁹:

$$\|\mathbf{a}\|_1 = \sum_{i=1}^n |\mathbf{a}[i]| \quad (28.13)$$

- The Euclidian norm:

$$\|\mathbf{a}\|_2 = \sqrt{\sum_{i=1}^n (\mathbf{a}[i])^2} \quad (28.14)$$

- The p -norm is a generalization of the two examples above:

$$\|\mathbf{a}\|_p = \left(\sum_{i=1}^n |\mathbf{a}[i]|^p \right)^{\frac{1}{p}} \quad (28.15)$$

- The infinity norm¹⁰ is the special case of the p -norm for $p \rightarrow \infty$:

$$\|\mathbf{a}\|_\infty = \max \{ |\mathbf{a}[1]|, |\mathbf{a}[2]|, \dots, |\mathbf{a}[n]| \} \quad (28.16)$$

Such norms can be used as distance measures, and we hence define a new distance measurement function as:

$$\text{dist}_{n,p}(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\|_p \quad \forall \mathbf{a}, \mathbf{b} \in A \subseteq \mathbb{R}^n \quad (28.17)$$

$$\text{dist}_{Man} \equiv \text{dist}_{n,1} \quad (28.18)$$

$$\text{dist}_{eucl} \equiv \text{dist}_{n,2} \quad (28.19)$$

If the places of the vectors $\mathbf{a}$ have different ranges, for example $\mathbf{a}[1] \in [0..1]$ and $\mathbf{a}[2] \in [0..100\,000]$, a norm of the difference of two such vectors may not represent their true "semantic" distance. Here, the contribution of the first elements of two vectors $\mathbf{a}$ and $\mathbf{b}$ to their distance will most likely be negligible. However, the two vectors $(0, 0)^T$ and $(1, 100)^T$ may be considered "more different" than $(0, 0)^T$ and $(0.1, 110)^T$, since they differ the whole range in their first elements. Therefore, an additional distance measure, the $\text{dist}_{n,p}^n$ distance is defined which normalizes the vector places before finally computing the norm.

$$\text{dist}_{n,p}^n(\mathbf{a}, \mathbf{b}) = \left(\sum_{i=1}^n \left| \begin{cases} \frac{\mathbf{a}[i] - \mathbf{b}[i]}{\text{range}(\mathbf{A})[i]} & \text{if } \text{range}(\mathbf{A})[i] > 0 \\ \mathbf{a}[i] - \mathbf{b}[i] & \text{otherwise} \end{cases} \right|^p \right)^{\frac{1}{p}} \quad (28.20)$$

⁶ http://en.wikipedia.org/wiki/Euclidean_distance [accessed 2007-07-03]

⁷ http://en.wikipedia.org/wiki/Pythagorean_theorem [accessed 2007-07-03]

⁸ http://en.wikipedia.org/wiki/Vector_norm [accessed 2007-07-03]

⁹ http://en.wikipedia.org/wiki/Taxicab_geometry [accessed 2007-07-03]

¹⁰ http://en.wikipedia.org/wiki/Maximum_norm [accessed 2007-07-03]

28.1.3 Elements Representing a Cluster

We already stated that there is not necessarily an $a \in A$ assigned to each real vector in $\mathbb{R}^n$. Thus, there also does not necessarily exist an element a in the center centroid c of a cluster c . For our purposes in this book, we are however interested in elements representing clusters. Since I have not found any other in literature, we will call such elements nuclei. We can define different functions $\text{nucleus}(c \in C)$ to compute such nuclei which, in turn, depend on a distance measure. We will again abbreviate this distance function by dist . dist is an implicit parameter which can be replaced by any of the functions introduced before. Also again, the default setting is $\text{dist} = \text{dist}_{eucl} \equiv \text{dist}_{n,2}$.

The first possible nucleus method, nucleus_c , would be to take the element which is closest to the centroid $\text{centroid}(c)$ of the cluster c :

$$n \in c = \text{nucleus}(c) \Leftrightarrow \text{dist}(a, \text{centroid}(c)) \geq \text{dist}(n, \text{centroid}(c)) \quad \forall a \in c \quad (28.21)$$

28.1.4 Distance Measures Between Clusters

In order to determine the distance between two clusters, another set of distance measures can be applied. Such distance measures usually will compute the distance between two clusters as a function of the distances between their elements which is, in turn, defined using a secondary distance function. We will abbreviate this secondary distance function by dist which can be replaced by any of the functions named in the above subsections. We assume it to be an implicit parameter with the default value $\text{dist} = \text{dist}_{eucl} \equiv \text{dist}_{n,2}$. Let c_1 and c_2 be two clusters in C , then we can define the following distance measures between them:

- The maximum distance between the elements of the two clusters (also called complete linkage):

$$\text{dist}_{max}(c_1, c_2) = \max \{ \text{dist}(a, b), \forall a \in c_1, b \in c_2 \}; \quad \forall c_1, c_2 \in C \quad (28.22)$$

- The minimum distance between the elements of the two clusters (also called single linkage):

$$\text{dist}_{min}(c_1, c_2) = \min \{ \text{dist}(a, b), \forall a \in c_1, b \in c_2 \}; \quad \forall c_1, c_2 \in C \quad (28.23)$$

- The mean distance between the elements of the two clusters (also called average linkage):

$$\text{dist}_{avg}(c_1, c_2) = \frac{1}{|c_1| * |c_2|} \sum_{\forall a \in c_1} \sum_{\forall b \in c_2} \text{dist}(a, b); \quad \forall c_1, c_2 \in C \quad (28.24)$$

- The increase in variance $\text{dist}_{var}(a, b)$ if the clusters were merged.
- The distance of their centers:

$$\text{dist}_{cent}(c_1, c_2) = \text{dist}(\text{centroid}(c_1), \text{centroid}(c_2)); \quad \forall c_1, c_2 \in C \quad (28.25)$$

- The distance of their nuclei computed by the nucleus function nucleus (see the definition of nucleus in Section 28.1.3):

$$\text{dist}_{nucl}(c_1, c_2) = \text{dist}(\text{nucleus}(c_1), \text{nucleus}(c_2)); \quad \forall c_1, c_2 \in C \quad (28.26)$$

28.2 Clustering Algorithms

28.2.1 Cluster Error

The most commonly used partitional clustering strategies are based on the square error criterion. The general aim is to obtain a partition which minimizes the square error for a given number k [1790] which we generalize to fit any given distance measure dist :

Definition 28.10 (Clustering Error). We define the error error_c inside a cluster as the sum of the distances of its elements from its center basing on a distance measure function. The total error of a partition error_p is then the sum of all the errors of the clusters included. Normally, we will use $\text{dist}_{\text{euc}} \equiv \text{dist}_{n,2}$ as distance measure.

$$\text{error}_c(c) = \sum_{\forall a \in c} \text{dist}(a, \text{centroid}(c)) \quad (28.27)$$

$$\text{error}_p(C) = \sum_{\forall c \in C} \text{error}_c(c) = \sum_{\forall c \in C} \sum_{\forall a \in c} \text{dist}(a, \text{centroid}(c)) \quad (28.28)$$

Normally, this error is minimized under the premise of a fixed number of clusters $k = |C|$. Then, an optimum configuration C^* is searched within the set $\mathfrak{C}_k$ of all such partitions of A into k clusters C . This optimum C^* is defined by $\text{error}_p(C^*) = \min \{\text{error}_p(C) \mid \forall C \in \mathfrak{C}_k\}$. Since testing all possible configurations C is too expensive (see Equation 28.7), finding the optimum C^* is an optimization task itself. Doing so inside an optimization process itself hence only rarely is applicable. Here we will introduce some algorithms which approximate good C .

28.2.2 k -means Clustering

k -means clustering¹¹ [1791, 1570, 1792] partitions the data points $a \in A$ into k disjoint subsets $c \subseteq A$, $c \in C \subseteq \mathfrak{C}_k$. It tries to minimize the sum of all distance of the data points and the centers of the clusters they belong to. In general, the algorithm does not achieve a global minimum of over the assignments. Despite this limitation, k -means clustering is used frequently as a result of its ease of implementation. [1793]

k -means clustering works approximately as follows [1790]:

1. Select an initial partition of k clusters.
2. Create a new partition by assigning each $a \in A$ to the cluster with the closest center. Repeat this until the partition does not change anymore.
3. Modify the cluster set by merging, dividing, deleting or creating cluster. If the clustering error of the new partition is smaller than the error of the previous one then go back to step 2.

In order to perform the modification of the cluster set, we introduce a function called `kMeansModify` obeying the following conditions.

$$\begin{aligned} C_{\text{new}} = \text{kMeansModify}_k(C) \Rightarrow & \forall a \in c_1 \in C \exists c_2 \in C_{\text{new}} : a \in c_2 \wedge \\ & \forall a \in c_2 \in C_{\text{new}} \exists c_1 \in C : a \in c_1 \end{aligned} \quad (28.29)$$

In other words, `kMeansModify` translates one set of clusters C to another one C_{new} by redeeming Definition 28.2 on page 468. One (crude) example for an implementation of `kMeansModify` is specified as Algorithm 28.1.

We demonstrate how k -means clustering works in Algorithm 28.2. As distance measure `dist` (lines 23 and 25) usually the Euclidian distance between the centroids of the clusters c_1 and c_2 , $\text{dist}_{\text{cent}}(c_1, c_2)$, see page 471, is used.

28.2.3 n^{th} Nearest Neighbor Clustering

The n^{th} nearest neighbor clustering algorithm is defined in the context of this book only. It creates at most k clusters where the first $k - 1$ clusters contain exactly one element. The remaining elements are all together included in the last cluster. The elements of the single-element clusters are those which have the longest distance to their n^{th} -nearest neighbor.

¹¹ http://en.wikipedia.org/wiki/K-nearest-neighbor_estimator [accessed 2007-07-03]

Algorithm 28.1: $C_{new} \leftarrow \text{kMeansModify}_k(C)$

Input: [implicit] k : the number of clusters wanted, $k \leq |A|$
Input: [implicit] dist : the distance measure between clusters to be used
Input: [implicit] dist_2 : the distance measure between elements to be used
Input: C : the list of clusters c to be modified
Data: m : the index of the cluster $C_{[m]}$ with the lowest error
Data: n : the index of the cluster $C_{[n]}$ nearest to $C_{[m]}$
Data: s : index of the cluster $C_{[s]}$ with the highest error
Output: C_{new} : the modified list of clusters

```

1 begin
2    $m \leftarrow m : \text{error}_c(C_{[m]}) = \min \{C_{[i]} \mid i \in [0, k-1]\}$ 
3    $n \leftarrow n : \text{dist}(C_{[m]}, C_{[n]}) = \min \{\text{dist}(C_{[m]}, C_{[i]}) \mid i \in [0, k-1] \setminus \{m\}\}$ 
4    $s \leftarrow s : \text{error}_c(C_{[s]}) = \max \{\text{error}_c(C_{[i]}) \mid i \in [0, k-1] \setminus \{m, n\}\}$ 
5    $C_{[m]} \leftarrow C_{[m]} \cup C_{[n]}$ 
6    $a \leftarrow a \in C_{[s]} : \text{dist}_2(a, \text{centroid}(C_{[s]})) \geq \text{dist}_2(b, \text{centroid}(C_{[s]})) \quad \forall b \in C_{[s]}$ 
7    $C_{[n]} \leftarrow \{a\}$ 
8    $C_{[s]} \leftarrow C_{[s]} \setminus \{a\}$ 
9   return  $B$ 
10 end
  
```

This clustering algorithm is suitable for reducing a large set to a smaller one which contains still the most interesting elements (those in the single-element clusters). It has relatively low complexity and thus runs fast, but on the other hand has the setback that dense aggregations of $\geq n$ elements will be put into the “rest elements”-cluster. For n , normally a value of $n = \sqrt{k}$ is used.

n^{th} nearest neighbor clustering uses the k^{th} nearest neighbor distance function $\text{dist}_{nn,k}^\rho$ introduced in Definition 27.62 on page 457 with its parameter k set to n . Do not mix this parameter up with the parameter k of this clustering method – although they have the same name, they are not the same. I know, I know, this is not pretty.

Notice that Algorithm 28.3 should only be applied if all the elements $a \in A$ are unique, i.e., there exists no two equal elements in A which is, per definition, true for all sets. In a real implementation, a preprocessing step should remove are duplicates from A before clustering is performed. Especially our home-made nearest neighbor clustering variant is unsuitable to process lists containing the same elements multiple times. Since all equal elements have the same distance to their n^{th} neighbor, it is likely that the result of the clustering is very unsatisfying since one element may occur multiple times whereas a variety of different other elements is ignored. Therefore, the aforementioned preprocessing should be applied, which may have the drawback that we could possibly obtain a set C with less than k clusters. In the Sigoa system’s implementation of the n^{th} nearest neighbor clustering, only one instance of each group of equal elements in A is permitted to become a single-node cluster per run and multiple runs are performed until k clusters have been created.

28.2.4 Linkage Clustering

The linkage method [1496, 1794] is used to create a set C containing at most k clusters. This algorithm initially creates a cluster of each single element in the set A . This set C of cluster c is reduced melting together the two closest clusters iteratively. Again, the distance measure function dist (see line 11 of Algorithm 28.4) used can be any of distance measures already introduced.

According to the cluster distance measure dist chosen, `linkageCluster` realizes different types of linkage clustering algorithms¹² (see Section 28.1.4 on page 471):

¹² http://en.wikipedia.org/wiki/Data_clustering#Agglomerative_hierarchical_clustering
 [accessed 2007-07-03]

Algorithm 28.2: $C \leftarrow \text{kMeansCluster}_k(A)$

Input: A : the set of elements a to be clustered
Input: [implicit] k : the number of clusters wanted, $0 < k \leq |A|$
Input: [implicit] dist : the distance measures between clusters to be used
Input: [implicit] dist_2 : the distance measures between elements to be used
Input: [implicit] kMeansModify : a function that modifies the cluster set
Data: C : the tuple of clusters c computed, $|C| = k$
Data: A_{cpy} : a temporary copy of A used for initialization
Data: C_{old} : the cluster set of the previous inner iteration
Data: C_{new} : the cluster set of the current inner iteration
Data: i : a counter variable for the loops
Data: d : the distance between the cluster $\{a\}$ and the current cluster in C_{old}
Data: $d_{\min}$: the minimum distance between $\{a\}$ and any cluster in C_{old}
Data: $i_{\min}$: the index of that cluster with the minimum distance in C_{old}
Output: c : the set of clusters – all the items of the tuple B represented as set

```

1 begin
2    $A_{\text{cpy}} \leftarrow A$ 
3    $k \leftarrow \min\{k, |A|\}$ 
4    $C_{\text{new}} \leftarrow \text{createList}(k, \emptyset)$ 
5    $i \leftarrow \text{len}(C_{\text{new}}) - 1$ 
6   while  $i > 0$  do
7      $C_{\text{new}}[i] \leftarrow \{a \in A_{\text{cpy}}\}$ 
8      $A_{\text{cpy}} \leftarrow A_{\text{cpy}} \setminus C[i]$ 
9      $i \leftarrow i - 1$ 
10   $C_{\text{new}}[0] \leftarrow A_{\text{cpy}}$ 
11  repeat
12     $C \leftarrow C_{\text{new}}$ 
13     $C_{\text{new}} \leftarrow \text{kMeansModify}_k(C_{\text{new}})$ 
14    repeat
15       $C_{\text{old}} \leftarrow C_{\text{new}}$ 
16       $i \leftarrow \text{len}(C_{\text{new}}) - 1$ 
17      while  $i > 0$  do
18         $C_{\text{new}}[i] \leftarrow \emptyset$ 
19         $i \leftarrow i - 1$ 
20      foreach  $a \in A$  do
21         $i \leftarrow \text{len}(C_{\text{old}}) - 1$ 
22         $i_{\min} \leftarrow 0$ 
23         $d_{\min} \leftarrow \text{dist}(\{a\}, C_{\text{old}}[0])$ 
24        while  $i > 0$  do
25           $d \leftarrow \text{dist}(\{a\}, C_{\text{old}}[i])$ 
26          if  $d < d_{\min}$  then
27             $d_{\min} \leftarrow d$ 
28             $i_{\min} \leftarrow i$ 
29           $i \leftarrow i - 1$ 
30         $C_{\text{new}}[i_{\min}] \leftarrow C_{\text{new}}[i_{\min}] \cup \{a\}$ 
31      until  $C_{\text{old}} = C_{\text{new}}$ 
32    until  $\text{error}_p(C) \leq \text{error}_p(C_{\text{new}})$ 
33    return listToSet( $C$ )
34 end
  
```

Algorithm 28.3: $C \leftarrow \text{nNearestNeighborCluster}_k(n) A$

Input: A : the set of elements a to be clustered
Input: [implicit] k : the number of clusters wanted ($0 < k \leq |A|$)
Input: [implicit] n : index for the nearest neighbors
Input: [implicit] dist : the distance measure to be used
Data: L : the sorted list of elements
Data: i : the counter variable
Output: C : the set of clusters c computed, $|C| = k$

```

1 begin
2    $L \leftarrow \text{sortList}_d(\text{setToList}(A), \text{dist}_{nn,k}^\rho \text{dist})$ 
3    $i \leftarrow \min\{k, |L|\} - 2$ 
4    $C \leftarrow \emptyset$ 
5   while  $i \geq 0$  do
6      $C \leftarrow C \cup \{L[i]\}$ 
7      $A \leftarrow A \setminus L[i]$ 
8      $i \leftarrow i - 1$ 
9   return  $C \cup \{A\}$ 
10 end

```

- If $\text{dist}(c_1, c_1) = \text{dist}_{\max}(c_1, c_1)$ denotes the maximum distance of the elements in two clusters, complete linkage clustering is performed.
- If $\text{dist}(c_1, c_1) = \text{dist}_{\text{avg}}(c_1, c_1)$ denotes the mean distance of the elements in two clusters, average linkage clustering is performed.
- If $\text{dist}(c_1, c_1) = \text{dist}_{\min}(c_1, c_1)$ denotes the minimum distance of the elements in two clusters, single linkage clustering is performed.

28.2.5 Leader Clustering

The leader clustering algorithm is a very simple one-pass method to create clusters. Basically, we begin with an empty leader list and an empty set of clusters. Step by step the elements a are extracted from the set A subject to clustering. a is then compared to the elements in the leader list in order to find one leader l with $\text{dist}(a, l)$ smaller than a specified maximum distance D . If such a leader exists, a is added to its cluster, otherwise a becomes leader of a new cluster containing only itself. The leader clustering can either be performed by using the first best leader l found with $\text{dist}(a, l) < D$ and assign a to its cluster ([619], Algorithm 28.5) or by comparing a to all possible leaders and thus finding the leader closest to a $\text{dist}(a, l) < \text{dist}(a, l_2) \forall l_2 \in \text{leaders}$ ([1795], Algorithm 28.6).

Algorithm 28.4: $C \leftarrow \text{linkageClusteribk}A$

Input: A : the set of elements a to be clustered
Input: [implicit] k : the number of clusters wanted ($0 < k \leq |A|$)
Input: [implicit] dist : the distance measure to be used
Input: [implicit] dist_2 : the distance measure between elements a to be used by dist
Data: c_1 : the first cluster to investigate
Data: c_2 : the second cluster to investigate
Data: d : the distance between the clusters r_1 and r_2 currently investigated
Data: $d_{\min}$: the minimum distance between two clusters c_{r_1}, c_{r_2} found in the current iteration
Data: c_{r_1} : the first cluster of the nearest cluster pair
Data: c_{r_2} : the second cluster of the nearest cluster pair
Output: C : the set of clusters c computed, $|C| = k$

```

1 begin
2    $C \leftarrow \emptyset$ 
3   foreach  $a \in A$  do  $C \leftarrow C \cup \{a\}$ 
4   while  $|C| > k$  do
5      $d_{\min} \leftarrow \infty$ 
6      $c_{r_1} \leftarrow \emptyset$ 
7      $c_{r_2} \leftarrow \emptyset$ 
8     foreach  $c_1 \in C$  do
9       foreach  $c_2 \in C$  do
10        if  $c_1 \neq c_2$  then
11           $d \leftarrow \text{dist}(c_1, c_2)$ 
12          if  $d \leq d_{\min}$  then
13             $d_{\min} \leftarrow d$ 
14             $b_{r_1} \leftarrow b_1$ 
15             $b_{r_2} \leftarrow b_2$ 
16       $C \leftarrow C \setminus c_{r_1}$ 
17       $C \leftarrow C \setminus c_{r_2}$ 
18       $C \leftarrow C \cup \{c_{r_1} \cup c_{r_2}\}$ 
19   return  $C$ 
20 end

```

Algorithm 28.5: $C \leftarrow \text{leaderCluster}_D^f(A)$ **Input:** A : the set of elements a to be clustered**Input:** [implicit] D : the maximum distance between an element and a cluster's leader**Input:** [implicit] dist : the distance measure to be used**Data:** a : an element in A **Data:** i : a counter variable**Data:** L : the list of cluster leaders**Output:** C : the list of clusters c computed

```

1 begin
2    $L \leftarrow ()$ 
3    $B \leftarrow ()$ 
4   foreach  $a \in A$  do
5      $i \leftarrow \text{len}(L) - 1$ 
6     while  $i \geq 0$  do
7       if  $\text{dist}(L[i], a) \leq D$  then
8          $C[i] \leftarrow C[i] \cup \{a\}$ 
9          $i \leftarrow -2$ 
10       $i \leftarrow i - 1$ 
11      if  $i \geq -1$  then
12         $L \leftarrow \text{addListItem}(L, a)$ 
13         $C \leftarrow \text{addListItem}(C, \{a\})$ 
14   return  $\text{listToSet}(C)$ 
15 end

```

Algorithm 28.6: $C \leftarrow \text{leaderCluster}_D^a(A)$ **Input:** A : the set of elements a to be clustered**Input:** [implicit] D : the maximum distance between an element and a cluster's leader**Input:** [implicit] dist : the distance measure to be used**Data:** a : an element in A **Data:** i : a counter variable**Data:** L : the list of cluster leaders**Output:** C : the list of clusters c computed

```

1 begin
2    $L \leftarrow ()$ 
3    $B \leftarrow ()$ 
4   foreach  $a \in A$  do
5      $i \leftarrow \text{len}(L) - 1$ 
6      $j \leftarrow 0$ 
7     while  $i > 0$  do
8       if  $\text{dist}(L[i], a) < \text{dist}(L[j], a)$  then  $j \leftarrow i$ 
9     if  $\text{dist}(L[j], a) \leq D$  then
10       $C[j] \leftarrow C[j] \cup \{a\}$ 
11    else
12       $L \leftarrow \text{addListItem}(L, a)$ 
13       $B \leftarrow \text{addListItem}(C, \{a\})$ 
14   return  $\text{listToSet}(C)$ 
15 end

```


Theoretical Computer Science

29.1 Introduction

Theoretical computer science¹ is the branch of computer science² that deals with the rather mathematical, logical, and abstract aspects of computing. It subsumes areas like algorithmic theory, complexity, the structure programming languages, and the solvability of problems.

29.1.1 Algorithms and Programs

In this and the following sections, we want to gain insight into the topic of algorithms, both in local and distributed systems. This seems to be appropriate, since any global optimization technique which we will discuss in this book is an algorithm. Often even a rather complicated one. Sometimes we even want to use several computers to solve an optimization problem cooperatively. Thus, we should know about the properties and theory of algorithms as well as of distributed systems.

The second reason is that many example applications discussed in this book will concern the automated syntheses of distributed algorithms. To understand these, knowledge of the features of distributed algorithms is valuable.

What are Algorithms?

The term *algorithm* comprises essentially all forms of “directives what to do to reach a certain goal”. A culinary receipt is an algorithm, for example, since it tells how much of what is to be added to the meal in which sequence and how it should be heated. The commands inside the algorithms can be very concise or very imprecise, depending on the area of application. How accurate can we, for instance, carry out the instruction “Add a tablespoon of sugar.”?

Hence, algorithms are such a wide field that there exist numerous different, rather fuzzy definitions for the word algorithm [1796, 1797, 1798, 1449, 1799]:

Definition 29.1 (algorithm). According to Whatis.com³, an algorithm is a procedure or formula for solving a problem. The word derives from the name of the mathematician, Mohammed ibn-Musa al-Khwarizmi, who was part of the royal court in Baghdad and who lived from about 780 to 850. Al-Khwarizmi’s work is the likely source for the word algebra as well.

¹ http://en.wikipedia.org/wiki/Theoretical_computer_science [accessed 2007-07-03]

² http://en.wikipedia.org/wiki/Computer_science [accessed 2007-07-03]

³ http://searchvb.techtarget.com/sDefinition/0,,sid8_gci211545,00.html [accessed 2007-07-03]

Definition 29.2 (algorithm). Wikipedia⁴ says that in mathematics, computing, linguistics, and related disciplines, an algorithm is a procedure (a finite set of well-defined instructions) for accomplishing some task which, given an initial state, will terminate in a defined end-state. The computational complexity and efficient implementation of the algorithm are important in computing, and this depends on suitable data structures.

Definition 29.3 (algorithm). An algorithm is a computable set of steps to achieve a desired result according to the National Institute of Standards and Technology⁵.

Definition 29.4 (algorithm). Wolfram MathWorld⁶ defines algorithm as a specific set of instructions for carrying out a procedure or solving a problem, usually with the requirement that the procedure terminate at some point. Specific algorithms sometimes also go by the name method, procedure, or technique. The word "algorithm" is a distortion of al-Khwarizmi, a Persian mathematician who wrote an influential treatise about algebraic methods. The process of applying an algorithm to an input to obtain an output is called a computation.




Fig. 29.1: The relation between algorithms and programs.

Whereas an algorithm is a set of directions in a representation that is especially understandable for human beings, programs are intended to be processed by machines and are therefore expressed in a more machine-friendly form. Originally, this was machine code. But for more than sixty years [1800], huge effort is being spent to allow us to write programs in

⁴ <http://en.wikipedia.org/wiki/Algorithm> [accessed 2007-07-03]

⁵ <http://www.nist.gov/dads/HTML/algorithm.html> [accessed 2007-07-03]

⁶ <http://mathworld.wolfram.com/Algorithm.html> [accessed 2007-07-03]

more and more comprehensible syntax. One could say that a program is basically an algorithm realized for a given computer or execution environment, as illustrated in Figure 29.1. The difference between programs and algorithms hence today lies primarily in the intention.

Definition 29.5 (Program). A program⁷ is a set of instructions that describe a task or an algorithm to be carried out on a computer. Therefore, the primitive instructions of the algorithm must be expressed either in a form that the machine can process directly (machine code⁸ [1801]), in a form that can be translated (1:1) into such code (assembly language [1802], Java byte code [1692], etc.), or in a high-level programming language⁹ [1803] which can be translated (n:m) into the latter using special software (compiler) [1804].

Definition 29.6 ((Software) Process). In terms of software, a process¹⁰ is a program that is currently executed. While a program only is a description of what to do, a process is the procedure of actually doing it. In a program for example the number and types of variables are described – in a process they are allocated and used.

Here we should also mention one of the most fundamental principle of electronic data processing¹¹, the IPO Model¹². As sketched in Figure 29.2, it consists of three parts:



Fig. 29.2: A process in the IPO model.

- The input (IPO) is an external information or stimulus that enters the system.
- The processing step (IPO) is the set of all actions taken upon/using the input. In terms of software, these actions are performed by a process which is the running instance of a program.
- The output (IPO) comprises the results of the computation (processing phase) that leave the system.

29.1.2 Properties of Algorithms

Besides these definitions, algorithms all share the following properties. Well, with few exceptions that we also will elaborate on.

Definition 29.7 (Abstraction). An algorithm describes the process of solving a problem on a certain level of abstraction which is determined by the elementary algorithms and elementary objects it uses and the applied formalism. One of the most important methods of abstraction is the definition and reuse of sub-algorithms.

Definition 29.8 (Discrete). A discrete algorithm works step-wise, i. e., is build on base of atomic executable instructions.

⁷ http://en.wikipedia.org/wiki/Computer_program [accessed 2007-07-03]

⁸ http://en.wikipedia.org/wiki/Machine_code [accessed 2007-07-04]

⁹ http://en.wikipedia.org/wiki/High-level_programming_language [accessed 2007-07-03]

¹⁰ http://en.wikipedia.org/wiki/Process_%28computing%29 [accessed 2007-07-03]

¹¹ http://en.wikipedia.org/wiki/Electronic_data_processing [accessed 2007-07-03]

¹² http://en.wikipedia.org/wiki/IPO_Model [accessed 2007-07-03]

Definition 29.9 (Finite). The definition of a (static) finite algorithm has a limited length. The sequence of instructions of static finite algorithms is thus finite. During its execution, a (dynamic) finite algorithm uses only a limited amount of memory to store its interim results.

Definition 29.10 (Termination). Each execution of an algorithm terminates after a finite number of steps and returns its results.

Definition 29.11 (Determinism). In each execution step of a deterministic algorithm, there exists at most one way to proceed. If no way to proceed exists, the algorithm has terminated.

Deterministic algorithms do not contain instructions that use random numbers in order to decide what to do or how to modify data. Most of the optimization techniques included in this book are randomized algorithms. They hence are not deterministic. We give an introduction into this matter in Definition 29.18 on page 484.

Definition 29.12 (Determined). An algorithm is determined if it always yields the same results (outputs) for the same inputs.

29.1.3 Complexity of Algorithms

For most problems, there exists more than one approach that will lead to a correct solution. In order to find out which one is the “best”, we need some sort of metrics which we can compare [1805, 1806].

The most important measures obtained by analyzing an algorithm¹³ are the time that it takes to produce the wanted outcome and the storage space needed for internal data [1449]. We call those the time complexity and the space complexity dimensions. The time-complexity denotes how many steps algorithms need until they return their results. The space complexity determines how much memory an algorithm consumes at most in one run. Of course, these measures depend on the input values passed to the algorithm. If we have an algorithm that should decide whether a given number is prime or not, the number of steps needed to find that out will differ if the inputs are 1 or $2^{32582657} - 1$. Therefore, for both dimensions, the best-case, average-case, and the worst-case complexity exist.

In order to compare the time and space requirements of algorithms, some approximative notations have been introduced [1807, 1097, 1804]. As we just have seen, the time and space requirements of an algorithm normally depend on the size of its inputs. We can describe this dependency as a function of this size. In real systems however, the knowledge of the exact dependency is not needed. If we, for example, know that sorting n data elements with the Quicksort algorithm¹⁴ [1808, 1705] takes in average something about $n \log_2 n$ steps, this is sufficient enough, even if the correct number is $2n \ln n \approx 1.39n \log_2 n$.

The Big-**O**-family notations introduced by Bachmann [1809] and made known by Landau [1810] allow us to group functions together that rise at approximately the same speed.

Definition 29.13 (Big-O notation). The big-**O**¹⁵ notation is a mathematical notation used to describe the asymptotical upper bound of functions.

$$f(x) \in \mathbf{O}(g(x)) \Leftrightarrow \exists x_0, m \in \mathbb{R} : m > 0 \wedge |f(x)| \leq m|g(x)| \forall x > x_0 \quad (29.1)$$

In other words, a function $f(x)$ is in **O** of another function $g(x)$ if and only if there exists a real number x_0 and a constant, positive factor m so that the absolute value of $f(x)$ is smaller (or equal) than m -times the absolute value of $g(x)$ for all x that are greater than x_0 .

¹³ http://en.wikipedia.org/wiki/Analysis_of_algorithms [accessed 2007-07-03]

¹⁴ <http://en.wikipedia.org/wiki/Quicksort> [accessed 2007-07-03]

¹⁵ http://en.wikipedia.org/wiki/Big_O_notation [accessed 2007-07-03]

Therefore, $x^3 + x^2 + x + 1 = f(x) \in \mathbf{O}(x^3)$ since for $m = 5$ and $x_0 = 2$ it holds that $5x^3 > x^3 + x^2 + x + 1 \forall x \geq 2$.

In terms of algorithmic complexity, we specify the amount of steps or memory an algorithm needs in dependency on the size of its inputs in the big- $\mathbf{O}$ notation. A discussion of this topic and some examples can be found in Table 29.1.

class	examples	description
$\mathbf{O}(1)$	$f_1(x) = 2^{222}$, $f_2(x) = \sin x$	Algorithms that have constant runtime for all inputs are $\mathbf{O}(1)$.
$\mathbf{O}(\log n)$	$f_3(x) = \log x$, $f_4(x) = f_4(\frac{x}{2}) + 1$; $f_4(x < 1) = 0$	Logarithmic complexity is often a feature of algorithms that run on binary trees or search algorithms in ordered sets. Notice that $\mathbf{O}(\log n)$ implies that only parts of the input of the algorithm is read/regarded, since the input has length n and we only perform $m \log n$ steps.
$\mathbf{O}(n)$	$f_5(x) = 23n + 4$, $f_6(x) = \frac{n}{2}$	Algorithms of $\mathbf{O}(n)$ require to access and process their input a constant number of times. This is for example the case when searching in a linked list.
$\mathbf{O}(n \log n)$	$f_7(x) = 23x + x \log 7x$	Many sorting algorithms like quicksort and mergesort are in $\mathbf{O}(n \log n)$
$\mathbf{O}(n^2)$	$f_8(x) = 34x^2$, $f_9(x) = \sum_{i=0}^{x+3} x - 2$	Some sorting algorithms like selection sort have this complexity. For many problems, $\mathbf{O}(n^2)$ -solutions are acceptable good.
$\mathbf{O}(n^i) : i > 1, i \in \mathbb{R}$	$f_{10}(x) = x^5 - x^2$	The general polynomial complexity. In this group we find many algorithms that work on graphs.
$\mathbf{O}(2^n)$	$f_{11}(x) = 23 * 2^x$	Algorithms with exponential complexity perform slowly and fast become unfeasible with increasing input size. For many hard problems, there exist only algorithms of this class. Their solution can otherwise only be <i>approximated</i> by the means of randomized global optimization techniques.

Table 29.1: Some examples of the big- $\mathbf{O}$ notation

Definition 29.14 (Big- Ω notation). The big- Ω notation is a mathematical notation used to describe the asymptotical lower bound of functions.

$$f(x) \in \Omega(g(x)) \Leftrightarrow \exists x_0, m \in \mathbb{R} : m > 0 \wedge |f(x)| \geq m|g(x)| \forall x > x_0 \quad (29.2)$$

$$f(x) \in \Omega(g(x)) \Leftrightarrow g(x) \in \mathbf{O}(f(x)) \quad (29.3)$$

Definition 29.15 (Θ notation). The Θ notation is a mathematical notation used to describe both, an upper and a lower asymptotical bound of functions.

$$f(x) \in \Theta(g(x)) \Leftrightarrow f(x) \in \mathbf{O}(g(x)) \wedge f(x) \in \Omega(g(x)) \quad (29.4)$$

Definition 29.16 (Small-o notation). The small-o notation is a mathematical notation used to define that a function is asymptotical negligible compared to another one.

$$f(x) \in \mathbf{o}(g(x)) \Leftrightarrow \lim_{n \rightarrow \infty} \left| \frac{f(x)}{g(x)} \right| = 0 \quad (29.5)$$

Definition 29.17 (Small- ω notation). The small- ω notation is a mathematical notation used to define that another function is asymptotical negligible compared to a special function.

$$f(x) \in \omega(gx) \Leftrightarrow \lim_{n \rightarrow \infty} \left| \frac{f(x)}{g(x)} \right| = \infty \quad (29.6)$$

$$f(x) \in \omega(g(x)) \Leftrightarrow g(x) \in \mathbf{o}(f(x)) \quad (29.7)$$

29.1.4 Randomized Algorithms

Deterministic algorithms¹⁶ will always produce the same results when given the same inputs. Such behavior comes closest to the original intention behind the definition of algorithms. The execution of a recipe should always yield the same meal, sorting identical lists should always result in, again identical, sorted lists. So in general, algorithms are considered to be deterministic. For many problems however, deterministic algorithms are unfeasible. In global optimization (see Section 1.1.1 on page 4), the problem space $\mathbb{X}$ is often extremely large and the relation of an element's structure and its utility as solution is not obvious. Hence, the search space $\mathbb{G}$ often cannot be partitioned wisely and an exhaustive search would be the only deterministic option left. Such an approach would take an infeasible long time. Here, the only way out is using a randomized algorithm.

Definition 29.18 (Randomized Algorithm). A randomized algorithm¹⁷ includes at least one instruction that acts on the basis of random numbers. In other words, a randomized algorithm violates the constraint of determinism. Randomized algorithms are also often called probabilistic algorithms [1811, 1812, 1813, 1814, 1815].

There are two general classes of randomized algorithms: Las Vegas and Monte Carlo algorithms.

Definition 29.19 (Las Vegas Algorithm). A Las Vegas algorithm¹⁸ is a randomized algorithm that never returns a wrong result [1797, 1811, 1812, 1814].

Either it returns the correct result, reports a failure, or does not return at all. If a Las Vegas algorithm returns, its outcome is deterministic (but not the algorithm itself). The termination (see Definition 29.10 on page 482) however cannot be *guaranteed*. There usually exists an *expected* runtime limit for such algorithms – their actual execution however may take arbitrarily long. In summary, we can say that a Las Vegas algorithm terminates with a positive probability and is (partially) correct.

Definition 29.20 (Monte Carlo Algorithm). A Monte Carlo algorithm¹⁹ always terminates. Its result however can be correct or incorrect [1811, 1812, 1814]. In contrast to Las Vegas algorithms, Monte Carlo algorithms always terminate but are (partially) correctly only with a positive probability.

Definition 29.21 (Monte Carlo Method). Monte Carlo methods²⁰ are a class of Monte Carlo algorithms used for simulating the behavior of systems of different types. Therefore, Monte Carlo methods are nondeterministic and often incorporate random numbers [1816, 1817, 1818, 1819].

¹⁶ http://en.wikipedia.org/wiki/Deterministic_computation [accessed 2007-07-03], see also Definition 29.11 on page 482

¹⁷ http://en.wikipedia.org/wiki/Randomized_algorithm [accessed 2007-07-03]

¹⁸ http://en.wikipedia.org/wiki/Las_Vegas_algorithm [accessed 2007-07-03]

¹⁹ http://en.wikipedia.org/wiki/Monte_carlo_algorithm [accessed 2007-07-03]

²⁰ http://en.wikipedia.org/wiki/Monte_Carlo_method [accessed 2007-07-03]

29.2 Distributed Systems and Distributed Algorithms

Definition 29.22 (Distributed System). A distributed system is a system of autonomous computers (nodes) which are connected loosely by a network and communicate by the exchange of messages in order to together perform a common functionality. Mühl [1820]

Definition 29.23 (Distributed Algorithm). Distributed algorithms [1821, 1462, 1461] are algorithms that are executed by multiple computers in such a distributed system cooperatively.

The different nodes in the network (and hence, the different instances of the distributed algorithm running on them) do not share the same view on the global state. They exchange information by the means of communication. Another source for differences in the view that the single nodes have on the global state may result from the lack of a globally synchronized time. This is due to the fact that communication involves usually latency²¹. If one node sends a message (including the current time, for instance) to another node, it takes some time t for the message to reach its destination. During this time, the message is regarded as *sent* by the first node and not yet known to the second one. When it arrives, its contents may already be outdated or it may have even been overtaken by another message which has been sent later. Furthermore, networks may induce arbitrary errors into the message's content and messages can even get lost.

Another aspect to consider when thinking about distributed algorithms is that the nodes in a network are not necessarily homogeneous. The computers in a distributed system may provide different computational power. This in turn will lead to different speed of progression of the single instances of the distributed algorithm.

Distributed algorithms can be distinguished from sequential algorithms because they run on multiple nodes in parallel in order to cooperatively solve one problem. They can be distinguished from mere parallel algorithms since node in the distributed system executes an instance of the same algorithm with a different view on the global state.

Distributed algorithms can provide the following advantages depending on their design:

1. modularity
2. flexibility
3. resource-sharing
4. no central point of failure because of decentralization
5. scalability because of decentralization
6. robustness
7. availability
8. fault-tolerance

Distributed algorithms may come with the following drawbacks, again depending on their design:

1. high complexity
2. no common view on the global state
3. no global time
4. processes may fail
5. latency in communication (see Section 29.2.3 on page 499)
6. faults in communication (Section 29.2.3 on page 498 and Section 29.2.3)
7. problems in termination detection
8. phantom/pseudo-deadlocks
9. race conditions

²¹ Which will be discussed in Section 29.2.3 on page 499.

The quality of a distributed algorithm can be determined by its communications complexity, i. e., how many messages need to be exchanged, or by its time complexity, i. e., how many computational steps need to be performed on the single nodes.

TODO

Definition 29.24 (Scalability). Scalability²² is a measure describing how good a system can grow or be extended for processing a higher computational load.

Definition 29.25 (Central Point Of Failure). A central (or single) point of failure is a subsystem or process that, if it fails, leads to the collapse of the whole distributed system. An example for central point of failures is central servers.

Definition 29.26 (Bottleneck). The bottleneck²³ of a distributed application is the part that has the most limiting influence on its performance.

Imagine, for instance, an hourglass. Here, the dilation in its center is the *bottleneck* that limits the amount of sand that can fall down per time unit.

TODO

29.2.1 Network Topologies

Definition 29.27 (Network Topology). Network topology²⁴ is the study of arrangement and mapping of the components of a network such as connections and nodes. One may also refer to a network arrangement itself as a topology, i. e., you can say “The topology of our network is a star.”

In the further text, we will use the term *edge* synonymously for link and connection and the term *vertex* as synonym for node or computer since network topology is closely related to graph theory.

Each computer network has exactly one physical topology which is the layout of its physical components (computers, cables). This physical structure defines which nodes can communicate directly with each other and which not. On top of that physical design, several virtual/overlay topologies may be built.

Definition 29.28 (Overlay Network). An overlay network²⁵ is a virtual network which is built on top of another computer network. The nodes in the overlay network are connected by virtual or logical links [1822].

IP addresses²⁶, for instance, form an overlay topology on top of MAC addresses²⁷ in Ethernets²⁸. A peer-to-peer network is an overlay network because it runs on top of the internet. Several distributed algorithms require the nodes to be arranged in special topologies like stars or rings. This can be achieved in arbitrary networks by defining an overlay structure which performs according routing and address translations.

When speaking of topology, one would normally think about a hardwired network of computers, connected with each other through Ethernet cabling and such and such. If we

²² <http://en.wikipedia.org/wiki/Scalability> [accessed 2008-02-08]

²³ <http://en.wikipedia.org/wiki/Bottleneck> [accessed 2007-07-03]

²⁴ http://en.wikipedia.org/wiki/Network_topology [accessed 2007-07-03]

²⁵ http://en.wikipedia.org/wiki/Overlay_network [accessed 2007-07-03]

²⁶ http://en.wikipedia.org/wiki/IP_address [accessed 2008-02-09]

²⁷ http://en.wikipedia.org/wiki/Mac_address [accessed 2008-02-09]

²⁸ <http://en.wikipedia.org/wiki/Ethernet> [accessed 2008-02-09]

consider a WLAN²⁹ or a wireless sensor network as described in Definition 29.32 on page 491 on the other hand, there is of course no such thing as cabling. But still, there is a certain topology: not all nodes may be able to directly contact each other since their radio transmission ranges are limited. They may only be able to exchange messages directly with some nodes in their physical neighborhood only. Hence, we can span a graph over this network, where each node is connected to his neighbors in communication range only. This graph then defines the topology.




Fig. 29.3.1: unrestricted topology



Fig. 29.3.2: bus




Fig. 29.3.3: star




Fig. 29.3.4: ring




Fig. 29.3.5: hierarchy




Fig. 29.3.6: grid




Fig. 29.3.7: fully connected

Fig. 29.3: Some simple network topologies

Unrestricted

In an unrestricted network topology as the one sketched in Fig. 29.3.1, we make only the general assumption that there is no network partition. In other words, for all nodes n in the network, there exists at least one path to each other node in the network. This path may, of course, consist of multiple hops over multiple connections.

²⁹ http://en.wikipedia.org/wiki/Wireless_LAN [accessed 2008-02-08]

Bus

All nodes in a bus system (illustrated in Fig. 29.3.2) are connected to the same transmission medium in a linear arrangement. All messages sent over the medium can be considered to be broadcasts that potentially can be received by all nodes more or less simultaneously. The transmission medium has exactly two ends.

Star

Fig. 29.3.3 shows an example for a star topology. Here, all nodes are connected to a single node in the center of the network. This center could, for example, be an Ethernet hub³⁰ or switch³¹ that retransmits the messages received to their correct destination. It could as well be a server that performs some specific tasks for the nodes. For a detailed discussion of the client-server architecture see Section 29.2.2.

Ring

In this topology, each node is connected to exactly two other nodes in a way that no partition exists. The nodes are arranged in a sequence where the first and the last node are connected with each other. An instance of the ring topology is illustrated in Fig. 29.3.4.

Hierarchy

Fig. 29.3.5 illustrates a hierarchical topology where the nodes of the network are arranged in form of a tree.

Grid

The nodes in a grid are laid out in a two-dimensional lattice so that each node (except those at the borders of the grid) is connected with four neighbors: one to the left, one to the right, one above and one below. Fig. 29.3.6 is an instance of such a topology.

Fully Connected

In a fully connected network, as sketched in Fig. 29.3.7, each node is directly connected with each other node.

29.2.2 Some Architectures of Distributed Systems

Client-Server Systems

Definition 29.29 (Client-Server). Client-server³² is a network architecture that separates two types of nodes: the client(s) and the server(s) [1823, 1824]. A client³³ utilizes a service provided by a server³⁴. It does so by sending a request to the server. This request contains details of the task to be carried out, for example the URL of a website to be returned. The server then executes appropriate actions and, in most cases, sends a response to the client. Usually, there is a small number of servers (normally one) which serves many clients.

³⁰ http://en.wikipedia.org/wiki/Ethernet_hub [accessed 2007-07-03]

³¹ http://en.wikipedia.org/wiki/Ethernet_switch [accessed 2007-07-03]

³² http://en.wikipedia.org/wiki/Client_server [accessed 2007-07-03]

³³ http://en.wikipedia.org/wiki/Client_%28computing%29 [accessed 2007-07-03]

³⁴ http://en.wikipedia.org/wiki/Server_%28computing%29 [accessed 2007-07-03]




Fig. 29.4: Multiple clients connected with one server

Client-server architectures like the one illustrated in Figure 29.4 are the most basic and the most common application logical architecture in distributed computing [1825, 1821, 1826]. They are part of almost all internet applications like:

- Websites³⁵ in the world wide web³⁶ are obtained by using the HTTP³⁷ protocol for communication between a web browser³⁸ and a web server³⁹.
- Application servers⁴⁰ contain the business logic of corporations. They support online shops⁴¹ with an underlying business model, for example.
- Database servers⁴² provide computers in a network with access to large data sets. Furthermore, they allow their clients to send structured queries that allow aggregation and selection of specific data.
- ...

The major advantages of client-server systems are their simplicity. Local algorithms can often be integrated into servers without too many problems while their adaptation to more complicated architectures is, well, more difficult and error-prone. The heaviest weakness of the client-server scheme is that the server represents a bottleneck (see Definition 29.25) and a single point of failure (see Definition 29.25 on page 486).

Peer-To-Peer Networks

Definition 29.30 (Peer-To-Peer Network). Instead of being composed of client and server nodes, a peer-to-peer⁴³ network consists only of equal peer nodes. A peer node works as a server for its fellow peers by providing certain functionality and simultaneously acts as client utilizing an similar service from its peers [1825, 1821, 1826, 1827, 1828]. Therefore, a peer node is often also called *servent*⁴⁴, a combination of the words server and client. The expression peer-to-peer is often abbreviated by P2P.

Peer-to-peer networks may have an arbitrary structure like the one sketched in Figure 29.5. While client-server systems are limited to providing communication between

³⁵ <http://en.wikipedia.org/wiki/Website> [accessed 2007-07-03]

³⁶ <http://en.wikipedia.org/wiki/WWW> [accessed 2007-07-03]

³⁷ <http://en.wikipedia.org/wiki/Http> [accessed 2007-07-03]

³⁸ http://en.wikipedia.org/wiki/Web_browser [accessed 2007-07-03]

³⁹ http://en.wikipedia.org/wiki/Web_server [accessed 2007-07-03]

⁴⁰ http://en.wikipedia.org/wiki/Application_server [accessed 2007-07-03]

⁴¹ http://en.wikipedia.org/wiki/Online_shop [accessed 2007-07-03]

⁴² http://en.wikipedia.org/wiki/Database_server [accessed 2007-07-03]

⁴³ <http://en.wikipedia.org/wiki/Peer-to-peer> [accessed 2007-07-03]

⁴⁴ <http://en.wikipedia.org/wiki/Servent> [accessed 2007-07-03]




Fig. 29.5: A peer-to-peer system in an unstructured network

the clients and the server solely, peer-to-peer networks may resemble any sort of underlying communication graph.

Peer-to-peer architectures circumvent the existence of single points of failure and can be constructed to be very robust against bottlenecks. They furthermore are often ad hoc, i. e., new peers may join the network at any time and leave it whenever they decide to do so. This can also be regarded as a drawback since the structure (and thus, its computational power and connectivity) of network may fluctuate heavily as well as the availability of data provided by the peers.

If obeying the definition exactly, there are no centralized components in a peer-to-peer network. There however exist hybrid networks where the peers for example register at a dedicated server which keeps track on the users online. Also, there exist different hierarchical or non-hierarchical overlay networks.

Important peer-to-peer-based applications are

- File and content sharing systems [1829, 1830] are the most influential and wide-spread P2P systems. Millions of users today share music, videos, documents and software over networks like Gnutella⁴⁵ [1631, 1632], Bittorrent⁴⁶, appleJuice⁴⁷ and the famous but shut-down Napster⁴⁸ network.
- Many scientific applications like Seti@home⁴⁹, Einstein@home⁵⁰, and Folding@home⁵¹ rely on users all over the world that voluntarily provide their unused computational resources. Such applications are most often constructed as screensavers that, after becoming active, download some pieces of data from a server and perform computations on them. After finishing the work on the received data, a response is issued to the server.
- Many instant messaging⁵² systems like talk⁵³ utilize peer-to-peer protocols. Most often, the clients need to log on and send status information to a server. Communication then

⁴⁵ <http://en.wikipedia.org/wiki/Gnutella> [accessed 2007-07-03]

⁴⁶ <http://en.wikipedia.org/wiki/BitTorrent> [accessed 2007-07-03]

⁴⁷ <http://www.applejuicenet.de/> [accessed 2007-07-03]

⁴⁸ <http://en.wikipedia.org/wiki/Napster> [accessed 2007-07-03]

⁴⁹ http://en.wikipedia.org/wiki/Seti_at_home [accessed 2007-07-03]

⁵⁰ <http://en.wikipedia.org/wiki/Einstein%40Home> [accessed 2007-07-03]

⁵¹ <http://en.wikipedia.org/wiki/Folding%40home> [accessed 2007-07-03]

⁵² http://en.wikipedia.org/wiki/Instant_messaging [accessed 2007-07-03]

⁵³ http://en.wikipedia.org/wiki/Talk_%28Unix%29 [accessed 2007-07-03]

either works client-server-based or in P2P-manner. Especially when audio or video chats come into play, peer-to-peer approaches are usually preferred.

- ...

Sensor Networks

Definition 29.31 (Sensor Network). A sensor network⁵⁴ [1667, 1831, 1832, 1833] is a network of autonomous devices which are equipped with sensors and together measure physical entities like temperature, sound, vibrations, pressure, motion, or such and such.

Definition 29.32 (Wireless Sensor Network). A wireless sensor network (WSN) [1834, 1668, 1835, 1836, 1837] is a sensor network where the single nodes are connected wirelessly, using techniques like wireless LAN⁵⁵, Bluetooth⁵⁶, or radio⁵⁷.




Fig. 29.6: A block diagram outlining building blocks of a sensor node.

Figure 29.6 sketches the building blocks of a sensor node. For communication with other nodes, short range radios, Bluetooth, or wireless LAN adapters are often added. The core of a sensor node is a microcontroller attached with RAM and ROM memory for code and data. The purpose of sensor networks is to measure some environmental parameters like temperature, humidity, or brightness. Thus, sensor nodes have one or multiple sensors attached.

Since they are autonomous devices usually not connected to power lines, sensor nodes have to be equipped with some sort of energy source. Chemical batteries are used to store energy, but often power scavenging units [1838, 1839, 1840, 1841] like, for example, solar cells [1842, 1843, 1844], thermal [1845] or kinetic energy harvesters [1846, 1847, 1848] are added. The field of energy supply of sensor nodes is critical and subject to active research [1849, 1850, 1851, 1852]. Batteries have limited capacity and are hard to replace after the network has been deployed. If no additional power scavenging unit is available, the sensor nodes will eventually stop functioning and become useless after all their energy is consumed. For

⁵⁴ http://en.wikipedia.org/wiki/Sensor_network [accessed 2007-07-03]

⁵⁵ http://en.wikipedia.org/wiki/Wireless_lan [accessed 2007-07-03]

⁵⁶ <http://en.wikipedia.org/wiki/Bluetooth> [accessed 2007-07-03]

⁵⁷ <http://en.wikipedia.org/wiki/Radio> [accessed 2007-07-03]

extending this lifetime, energy intense operations like communication via radio transmissions need to be reduced as much as possible.

The size of the sensor nodes ranges from shoe box down to matchbox dimensions. There is a strong wish to produce smaller and smaller nodes. Small sensors are recognized less obviously and blend better in their environment. Since they require fewer raw material, they might become much cheaper than their larger pendants. On the other hand, with this movement in the direction of sensor that are really tiny, some hard constraints arise. The size of the battery limits the amount of energy that can be stored, as well as the extent of a solar cell limits its energy production. The node size also restricts the dimensions of the memory and the sensors of the node [1853].

Other important research topics are data fusion and transportation in a WSN [1854, 1855] as deployment and maintenance [1856, 1857].

Widespread sensor node architectures are:

- *BTNodes*⁵⁸ are autonomous wireless communication and computing platforms based on a Bluetooth radio and a microcontroller. Developed at the ETH Zurich, BTNodes serve especially as demonstration, teaching, and research platforms. Fig. 29.7.1 shows a BTNode.
- Crossbow's *MICA2*⁵⁹ motes are multipurpose nodes. These systems are applied widely in real-world applications like environmental control in agriculture and outdoor sports as well as for indoor sports and military purposes. A picture of the Mica2Dot platform can be found in Fig. 29.7.2.
- Scatterweb⁶⁰ provide both, a research platform (*MSB* nodes, illustrated in Fig. 29.7.3) and an industrial sensor network (*ScatterNodes*).
- Dust Networks⁶¹ developed their *SmartMesh* for building wireless solutions for the global market. Their nodes use the Time Synchronized Mesh Protocol and middle-range radio to provide the reliability of a typical WLAN in their sensor networks. Fig. 29.7.4 shows a Dust Networks Evaluation Mote.
- ...

A small example application demonstrating the use of sensor networks is discussed in Section 23.2.2 on page 366.

Properties of Peer-To-Peer Systems and Sensor Networks

- Current peer-to-peer networks are often large-scale, with tens of thousands [1830] up to millions [1858] of users/nodes online. Although networks of thousands of sensors are a future goal, the number of nodes in sensor networks has not yet reached this extent. However, systems of several hundreds of nodes have already been deployed [1859, 1860].
- Since wireless sensor networks have limited transmission range, it is possible that not all nodes in a network can communicate directly with each other. The same issue exists in the internet, but there it is solved in a transparent manner by routers. In sensor networks however, dedicated hardware routers normally do not exist. Therefore, special routing protocols [1861, 1862, 1863] are applied. Here we see a strong relation between sensor networks and peer-to-peer systems: Each sensor may act as a sender of messages as well as router for other nodes. There is no generic hierarchy or division between senders or routers.
- Especially in peer-to-peer applications, there are strong fluctuations in the network membership. In content sharing networks for example, new users continuously join and leave

⁵⁸ <http://www.btnode.ethz.ch/> [accessed 2007-07-03]

⁵⁹ <http://www.xbow.com/Products/productdetails.aspx?sid=156> [accessed 2007-07-03]

⁶⁰ http://www.inf.fu-berlin.de/inst/ag-tech/scatterweb_net/ [accessed 2007-07-03] and <http://www.scatterweb.com/> [accessed 2007-07-03]

⁶¹ <http://www.dustnetworks.com/> [accessed 2007-07-03]




Fig. 29.7.1: BTNode




Fig. 29.7.2: Mica2Dot




Fig. 29.7.3: MSB Mote



Fig. 29.7.4: Dust Networks Evaluation Mote

Fig. 29.7: Images of some sensor network platforms

the network. In sensor networks on the other hand, volatility in the network structure arises from newly deployed nodes or nodes that become inactive because they ran out of battery power. A sensor node spends much of its time in sleep mode (so do I) and may be regarded as inactive in this time. When it triggers back to active mode, it again becomes member of the network. Furthermore, networks of mobile sensors have large fluctuations in their topology per default.

- Since sensor networks utilize sleep cycles in order to reduce energy consumption, messages that are routed may arbitrarily be delayed or even get lost.
- P2P networks often represent very heterogeneous environments, consisting of computers of different architectures and operating systems. Sensor networks on the other hand are most often homogeneous systems.

29.2.3 Modeling Distributed Systems

In this section we will discuss how to model the features of distributed systems. Models are needed to verify, test, and even prove the properties of distributed algorithms and are the foundation for any form simulation.

Modeling Parallelism

In reality, the distributed algorithms on the nodes of a network all run in parallel in an unpredictable manner. We now need to find a model that allows us to explore this parallel behavior in a formalized fashion.

In simulations, it is complicated to model time continuously as it would be physically correct. A discretization of time simplifies many aspects of a parallel system. We can justified

such an approach with the fact that time, as a continuum, can be split into infinite small units. We proceed analogously to our elaborations on the Poisson process in Section 27.3.2 on page 433. For time steps Δt of a short length, we make the assumptions that in such a time step either

1. nothing happens or that
2. exactly one node of the modeled systems performs exactly one elementary action.

With this simplification, we also indirectly have sequenced all elementary/atomic actions in our model of the running system. Furthermore, if the data exchange of the distributed algorithms can be modeled as instantaneous with discrete time, we can disregard the time steps where nothing happens in our considerations.

Each distributed algorithm to be analyzed now needs to be broken up into atomic actions. The way we do this has severe impact on the complexity of the analysis. Let us take Algorithm 29.1 for example, a very simple and naïve graph coloring algorithm. The graph coloring problem is defined on unrestricted graphs of arbitrary structure. The objective is to find a coloring where no node has the same color as any other node it is directly connected with. Our sample algorithm tries to solve this by first selecting a random color for the node it runs on and broadcasting this color to each of its neighbors. If a node now receives a message (containing the color of one of its neighbors), it checks if it has the same color set. If so, it randomly selects another one and again, informs its neighbors about its choice. From the looks, we could divide this algorithm into at least eight elementary decisions or actions. Since these consist of high-level constructs like *if* or *while*, we probably would have to further split them up and transform them into jumps in some sort of pseudo-assembler that we can simulate properly. Making formal statements about such aggregations of instructions in the context of parallelism gets more and more cumbersome the more instructions have to be taken into account.

Algorithm 29.1: distributedGraphColoring

Input: *colors*: a set of colors to choose from

Data: *color*: the color of the node

Data: *msg*: a message received from another node, containing the other node's color

```

1 begin
2   color  $\leftarrow$  colors[randomu(0, len(colors))]
3   broadcast(color)
4   while true do
5     msg  $\leftarrow$  receive()
6     if msg = color then
7       while msg = color do color  $\leftarrow$  colors[randomu(0, len(color))]
8       broadcast(color)
9 end
```

Reasoning about the features of our sample algorithm becomes way simpler when recognizing that we only need to decompose it into two pieces. The initialization part (lines 2 and 3) and the update part (lines 5 to 8). We can do this because only the sequence of those algorithm steps where a node receives or sends a message do matter in global sequence.⁶² Everything in between is local to the single node and has no influence on the progress of the algorithm instances on the other nodes. Therefore, we can assume the message sending/receiving actions together with all the actions in between them as elementary and atomic from the global standpoint.

⁶² receive() blocks and broadcast is non-blocking.

Reducing the atomic pieces of a distributed algorithm also reduces the number of relations to be taken into consideration and thus, eases its analysis significantly. The same goes for its simulation, since it is much more cumbersome to build interpreters for assembler-style elementary instructions than simple simulating a few statements in form of compound, high-level expressions.

After splitting up the algorithms into their elementary building blocks, we can classify parallelism models according to their degree of concurrency.

Synchronous Models

In synchronously running distributed systems, all nodes proceed with exactly the same speed. We model this by executing elementary algorithm steps according to the round robin principle⁶³. From the view of a single node, all its neighbors are always as fast as it is itself. This becomes clear when remember that the time where a node “does nothing”, i. e., where none of its atomic instructions is executed, does not exist from the viewpoint of the node. From the global perspective, runtime is assigned to the nodes as illustrated in Figure 29.8.




Fig. 29.8: Synchronous parallelism in a model of a network of five nodes.

Asynchronous Models

In some cases it may be sufficient to consider a distributed system as synchronous – most real networks however are not. They consist of computers that run very asynchronously, maybe even at different speeds that can change by and by. Such behavior can be modeled in a surprisingly simple fashion. Instead of assigning the runtime to the nodes using to the round robin principle, in each time step one node is picked randomly according to the uniform distribution (see Section 27.3.1 on page 430). Thus, in each time step each node has exactly the same probability of executing one action. In average over infinite time, all nodes will execute the same amount of algorithm steps. For a given time period however, it is possible that one node can execute five steps while another one just proceeds by two, as outlined in Figure 29.9. It is obvious that in this time period, node one progresses with more than double the speed than node four. On the other hand, during the following twenty time units, this ratio may as well be exactly reverted. We hence can model different and changing execution speeds. All nodes still have the same average speed, as it would be in a homogeneous network of computers of the same type. In case this is not wanted, we simply need to adjust the probability with which the nodes are picked and deviate from the uniform distribution.

The single nodes start at different points in time with their execution. Hence, one node can be done with its work before another one gets assigned the first time step. If the distributed algorithms do furthermore not contain infinite loops or similar constructs and thus, terminate at some point of time, this will result in a dynamic network topology where nodes join and leave the computation in an arbitrary manner.

⁶³ http://en.wikipedia.org/wiki/Round-robin_scheduling [accessed 2007-07-03]

nodes that are granted the same runtime, are modeled/simulated. Since the time units are assigned randomly to them, node 1 has spent up all his ticks before nodes 4 to 6 did even start. The network hence first consisted only of the node 1, then later of the nodes 1 to 3 and transcends over consisting of the nodes 4 to 6 to its final state where only the node 6 remains. Dynamic topologies should, in general, be regarded as unrestricted topologies. For




Fig. 29.10: Dynamic topology due to overlapping active times of nodes.

the ease of simulation implementation, we can however define a grid topology as basis where nodes are considered as switched on and off according to the random consumption of their runtime.

Static and Dynamic Partitions

Definition 29.33 (Static Partition). A static partition in a network $N = (n_1, n_2, \dots)$ containing the nodes n_i exists if we can divide the network in at least two disjoint subsets $N_1 \cup N_2 = N, N_1 \cap N_2 = \emptyset$ in way that there does not exist a connection between any node in N_1 and a node in N_2 .

If a network is statically partitioned, there exist at least two nodes that, even with routing over arbitrary many stations, will never be able to exchange data in any way. Such configurations are not interesting for simulations and models, since distributed algorithms cannot work on them properly. The modeler or simulation designer thus has to ensure that network partitions do not occur.

Communication

Every distributed algorithm requires some form of communication.

Number of Receivers

In a network, a node may be connected to n other nodes. Communication forms can be divided into three categories according to how many receivers are reached with one transmission as sketched in Figure 29.11.

Definition 29.34 (Unicast). Unicast⁶⁴ means sending a message to one single recipient only, as illustrated in Fig. 29.11.1.

Definition 29.35 (Multicast). Fig. 29.11.2 outlines the multicast⁶⁵ transmission scheme. A multicast message is sent to a subset of m recipients (in a group of n possible receivers), where $1 \leq m \leq n$.

⁶⁴ <http://en.wikipedia.org/wiki/Unicast> [accessed 2007-07-03]

⁶⁵ <http://en.wikipedia.org/wiki/Multicast> [accessed 2007-07-03]




Fig. 29.11.1: Unicast




Fig. 29.11.2: Multicast




Fig. 29.11.3: Broadcast

Fig. 29.11: The three different message transmission types.

Definition 29.36 (Broadcast). To broadcast⁶⁶ a message means to send it to all n possible recipients at once. Fig. 29.11.3 sketches a broadcast message that is received by all destinations that are directly linked with the sender.

If the group of possible recipients only has $n \leq 1$ members, there is no distinction between unicast, multicast, or broadcast. In general, we could regard broadcasts as a special case of multicast that is sent to all possible destinations. Furthermore, one could model unicasts as multicasts with only one recipient.

Reliability of Transmissions

The reliability of transmissions is another aspect that should be thought of when modeling the environment of a distributed algorithm. We can distinguish between two different types of faults in communication, message loss and message modification. In addition, there is the message delay/latency which can also take on the characteristics of an error.

When modeling a network in order to derive properties of a distributed algorithm, we have to decide which of these faults are relevant and which are not. If we, for example, examine an algorithm or program that runs in an environment where communication is secured by underlying protocol levels, message loss may be irrelevant. Also, message modification can be omitted from the model if the algorithms are not security-related and are expected to run on top of a protocol that uses checksums and such and such to secure data integrity later in the real implementation.

If the model bases on broadcast or multicast, we also have to decide whether possible faults occur per message or per packet on a connection. As already discussed before, we can think of a network as a graph where each node is represented by a vertex. A node is connected to all n nodes that it can directly communicate with by an edge. A broadcast message can be modeled by copying it n times. Each of these packages then travels along one of the connections. Here we should distinguish if we model a hard-wired network, where errors would occur packet-wise since they are usually bound to a single connection, or if we have a wireless network where an error would probably influence the whole transmission. The following considerations may be applied in either way.

Message Loss Messages may vanish on the way to their recipient. There are several possible reasons for such an error:

- This is most often caused by a collision⁶⁷ with another transmission on the same medium.
- The sending node may not be directly connected to the receiving one. Then, its messages have to be routed over some intermediate nodes. If one of these stations gets congested and the in- or output buffers of it are overflowing, it may discard the message.

⁶⁶ http://en.wikipedia.org/wiki/Broadcasting_%28networks%29 [accessed 2007-07-03]

⁶⁷ http://en.wikipedia.org/wiki/Collision_%28telecommunications%29 [accessed 2007-07-03]

- Messages can also be caused to disappear by a third party that intrudes the communication channel with the malicious intent of depriving the receiver of information.
- A connection between two nodes may be broken and all messages on this connection will be lost.

The loss of a message can be detected on the receiving side if the communication source sending a strictly increasing sequence number along with the payload, for instance. Well known protocols like TCP⁶⁸ utilize this mechanism.

If we need to model message loss, we can do so by simple determining a probability $0 \leq \varepsilon \leq 1$ with which transmissions fail. Then, whenever a message is sent by node, we use draw a random number e uniformly distributed in $[0, 1)$. If $e < \varepsilon$ then the transmission is lost, otherwise it will get through. This is a rather crude approach but has the advantage that it can easily be understood and analyzed using mathematical methods.

In order to add the possibility of a complete connection break down, we can draw another random number at each time step for a connection in order to determine whether it will fail or not. A failing connection then will stay broken for a time determined using the exponential distribution.

Message Modification There are two possible causes why messages could be modified in transmission.

- Accidental modification due to physical interference like power surges or other signals.
- Again, a third party may be responsible for the change of the message's content. The intention could be to deceive the receiver.

If we can assume that the distributed algorithms can rely on an underlying protocol with error detection capabilities, like the IP⁶⁹ which therefore uses a checksum⁷⁰, message modification does not need to be modeled.

Otherwise, we can regard a message m as a sequence of n bits $m = (b_1, b_2, \dots, b_n)$.

Definition 29.37 (Error Burst). An error burst⁷¹ is a very common error scheme in telecommunications. It denotes a continuous sequence of symbols (in our case, bits) over a data transmission channel or part of a message such that this sequence contains not a single correct (i. e., unaltered) symbol.

We can simulate such errors with Algorithm 29.2 by first defining a fault probability $0 \leq \varepsilon \leq 1$. For each message we draw a uniformly distributed random number from $[0, 1)$. If this number is smaller than ε , the message will be modified. Therefore, we first randomly determine the length of the error burst in line 4 and then its position. In Algorithm 29.2. We simply draw the error burst length from a normally distributed random variable with mean μ and variance σ^2 which have to be chosen carefully. In order to limit the burst length to a positive, natural number smaller or equal to the message length, we apply the cut-off function which has been introduced in Section 27.7.2 on page 462 and round down.

Message Latency

Definition 29.38 (Latency). Latency⁷² is the time difference between the moment where something is initiated and the moment when its effects becoming observable [1864].

⁶⁸ http://en.wikipedia.org/wiki/Transmission_Control_Protocol [accessed 2007-07-03]

⁶⁹ http://en.wikipedia.org/wiki/Internet_Protocol [accessed 2007-07-03]

⁷⁰ <http://en.wikipedia.org/wiki/Checksum> [accessed 2007-07-03]

⁷¹ http://en.wikipedia.org/wiki/Error_burst [accessed 2007-07-03]

⁷² http://en.wikipedia.org/wiki/Network_Latency [accessed 2007-07-03], <http://en.wikipedia.org/wiki/Lag> [accessed 2007-07-03]

Algorithm 29.2: $m_\varepsilon \leftarrow \text{errorBurst}(m, \varepsilon)$

Input: m : the original message
Data: l : the random length of the error burst
Data: p : the position where the error burst occurs
Output: m_ε : the message after passing the transmission channel

```

1 begin
2    $m_\varepsilon \leftarrow m$ 
3   if  $\text{random}_u() < \varepsilon$  then
4      $l \leftarrow \lfloor \text{random}_l(\text{random}_n(\mu, \sigma^2), 1, \text{len}(m) + 1) \rfloor$ 
5      $p \leftarrow \lfloor \text{random}_u(1, \text{len}(m) - l + 1) \rfloor$ 
6     while  $l > 0$  do
7        $m_\varepsilon[p] \leftarrow \overline{m[p]}$ 
8        $l \leftarrow l - 1$ 
9        $p \leftarrow p - 1$ 
10    return  $m_\varepsilon$ 
11 end
```

In a real distributed application, messages are constructed and sent by a software process running on a node. We consider the moment when they are passed down to the operating system or to the middleware as the moment when they are sent. From there, however, they have to be passed to the communication hardware from where they are transmitted over a medium. Now physical effects will delay the message from arriving instantly at the receiver side (the speed of light still cannot efficiently be surpassed). This delay induced by physical laws is normally negligible. Yet it is observable in satellite communication, for example when the host of a news show talks with a reporter on the other side of the globe. Besides electromagnetic transmission, there also other communication mechanisms where there is a noticeable delay induced by the medium. In biologically-inspired networks, for instance, data is transported by hormones or calcium molecules which certainly travel much slower than electromagnetic waves.

If the node sending a message is not directly connected to the recipient, the message will be routed over some intermediate nodes. Each of these nodes needs to analyze the message's destination in order to find out where to send it next which needs some processing time. Furthermore, congestion may lead to additional delay. After the message arrives on its destination, it has again to be dealt with by the hardware, operating system, and middleware before being passed to the application, which marks the arrival time.

The time a message needs to travel will of course also depend on its size. The most problematic fact is that there is no upper limit for message latency, i.e., we cannot determine if a message was lost or is still delayed in transmission.

Since the latency in a real network depends on so many factors [1864], it is hard to simply approximate it with a simple probability distribution. In principle, we can choose between three candidates: the exponential, the normal, and the uniform distribution. Additionally, latency can also be modeled deterministically as a function of the distance between the sender and the receiver [1865].

- Many experiments indicate that network latency is loosely exponentially distributed. On page three in [1866], the latency/probability diagram remotely resembles the exponential distribution. It is therefore often used in models and simulations [1867]. Here, a positive cut-off should be applied to prevent messages from traveling with zero delay (see Section 27.7.2 on page 462).
- Since the exact form of the distribution of latency may vary from application to application, it also makes sense to approximate it with a normal distribution. A normally distributed random number can take on zero or negative values with a non-zero proba-

bility. Therefore, again a cut-off should be applied in the model or simulation according to Section 27.7.2.

- Drawing uniformly distributed random numbers is the most crude and simple way to determine message latency. One reason for choosing it is that in most models, the exact distribution of the message delays plays no role. The only thing that counts is that messages can be delayed in a way that allows them to outpace each other or to allow some nodes to receive them early than others. These are the most critical scenarios for distributed algorithms and we can create them with the uniform distribution as well as with every other one.
- Another attractive approach to modeling latency is to make it a function of the distance between the sender and the receiver [1865]. If we consider a routed network and assume a constant delay per router passed by the message, this method is very elegant. Additionally, one could allow a bit of deviation by adding a small random number to each deterministically computed delay.

Reliability of Nodes

So far we have considered errors that may happen during message transmission. Now we want to discuss the errors that may happen before – the possible faults on the single nodes. Especially important in the context of distributed systems are Byzantine faults.

Definition 29.39 (Byzantine Fault). Byzantine faults⁷³ are errors that occur during the execution of a distributed algorithm in a network due to one or multiple nodes deviating from the prescribed flow of the algorithm [1868, 1869, 1870].

Constructing systems that are able to deal with such problems is subject to research since the 1980s [1871, 1872, 1873, 1874, 1875, 1876, 1877, 1878].

Byzantine faults cover node crashes, incorrect execution of algorithm steps, and the execution of wrong steps as well as nodes that *intentionally* send incorrect and misleading messages. Including Byzantine faults in the model of a distributed system in principle only needed in the last case, since the effects of the others can as well be reached by modeling communication faults.

29.3 Grammars and Languages

Languages are the most important means for communication between higher animals⁷⁴. Formal languages can also be used to define the formats for data being stored by or exchanged between computers and/or human beings. When analyzing a statement in a given language, we distinguish between its syntax and semantic.

Definition 29.40 (Syntax). The syntax⁷⁵ of a language is the set of rules that governs its *structure*. Each valid statement of a language must obey its syntactical structure. The sentence “*I am reading a book.*” is a sequence of a subject, a predicate, and an object.

Definition 29.41 (Semantic). The semantic⁷⁶ refers to the *meaning* of a statement. The sentence “*I am reading a book.*” has the meaning that the writer of it is visually obtaining information from a set of bounded pages filled with written words.

⁷³ http://en.wikipedia.org/wiki/Byzantine_fault_tolerance [accessed 2007-07-03]

⁷⁴ <http://en.wikipedia.org/wiki/Language> [accessed 2007-07-04]

⁷⁵ <http://en.wikipedia.org/wiki/Syntax> [accessed 2007-07-03]

⁷⁶ <http://en.wikipedia.org/wiki/Semantics> [accessed 2007-07-03]

29.3.1 Syntax and Formal Languages

Let us now take a closer look on the syntax of formal languages [1879, 1880].

Definition 29.42 (Alphabet). A finite set Σ of symbols (characters) $\alpha \in \Sigma$ with a total order (see Section 26.7.2 on page 415) defined on it is called an alphabet.

Definition 29.43 (Character String). A character string⁷⁷ (or word) over an alphabet Σ is any finite sequence of symbols $\alpha \in \Sigma$. Character strings have the following properties:

1. The empty character string ε is a character string over Σ .
2. If x is a character string over Σ , then αx is also a character string over Σ for all $\alpha \in \Sigma$.
3. β is a character string over the alphabet Σ if and only if it can be created using the two rules above.

Definition 29.44 (Concatenation). The concatenation⁷⁸ $\alpha\beta$ of two character strings $\alpha = \alpha_1\alpha_2\alpha_3\dots\alpha_n$ and $\beta = \beta_1\beta_2\beta_3\dots\beta_m$ over the alphabet Σ is the character string $\alpha\beta = \alpha_1\alpha_2\alpha_3\dots\alpha_n\beta_1\beta_2\beta_3\dots\beta_m$ which begins with α immediately followed (and ended by) β .

The set of all strings of length l over the alphabet Σ is called Σ^l with $\Sigma^0 = \{\varepsilon\} \forall \Sigma$. The set of all strings on Σ is called Σ^* , i. e., $\Sigma^* = \cup_{l=0}^{\infty} \Sigma^l$. It is also called Kleene star⁷⁹ (or Kleene closure).

Definition 29.45 (Lexeme). A lexeme⁸⁰ is the lowest level of syntactical unit of a language [1879]. It denotes a set of words that have the same meaning, like *run*, *runs*, *ran*, and *running* in English. A lexeme belongs to a particular syntactical category and has a semantic meaning.

Based on these definitions, we can consider a sentence to be a sequence of lexemes which, in turn, are character strings over some alphabet.

Definition 29.46 (Language). A language $\mathbb{L}$ over the alphabet Σ is a subset of Σ^* [1880]. $\mathbb{L}$ is the set of all sentences over an alphabet Σ that are valid according to its rules in syntax (the grammar) [1881].

When describing the formal syntax of a language, there are two possible approaches:

1. We can define recognizers that determine the structure of a sentence and can decide whether it belongs to the language or not. Recognizers are, for instance, used in compilers [1882].
2. A generative grammar can be defined from which all possible sentences of a language can be constructed.

29.3.2 Generative Grammars

A generative grammar G of a language $\mathbb{L}$ is a formal specification that allows us to construct every single sentence in $\mathbb{L}$ by applying recursive replacement rules. Therefore, we define non-terminal symbols (also called variables) which do not occur in the language's text and terminal symbols that do. One example of such a grammar is:

```

1 sentence  $\rightarrow$  subject verb object
2 subject  $\rightarrow$  Alice  $\vee$  Bob
3 verb  $\rightarrow$  writes  $\vee$  reads
4 object  $\rightarrow$  cipher-text  $\vee$  plain text

```

Listing 29.1: A simple generative grammar.

⁷⁷ http://en.wikipedia.org/wiki/Character_string [accessed 2007-07-03]

⁷⁸ <http://en.wikipedia.org/wiki/Concatenation> [accessed 2007-07-10]

⁷⁹ http://en.wikipedia.org/wiki/Kleene_star [accessed 2007-07-03]

⁸⁰ <http://en.wikipedia.org/wiki/Lexeme> [accessed 2007-07-03]

Here we have four productions, the terminal symbols **Alice**, **Bob**, **writes**, **reads**, **cipher-text**, **plain-text**, plus five non-terminal symbols (**sentence**, **subject**, **verb**, and **object**).

Definition 29.47 (Formal Grammar). A formal grammar⁸¹ $G = (N, \Sigma, P, S)$ is a 4-tuple consisting of:

- a finite set N of non-terminal symbols (variables),
- the alphabet Σ , a finite set of terminal symbols,
- a finite set P of productions (also called rules), and
- at least one start symbol $S \in N$ which belongs to the set of non-terminal symbols N .

Notice that the alphabet Σ here is not limited to letters or numerals, but may contain words, sentences, or even arbitrarily long texts. Additionally, we call the set $V = N \cup \Sigma$ including terminal and non-terminal symbols the grammar symbols.

The Chomsky Hierarchy

The Chomsky hierarchy stands for a hierarchy of formal grammars that generate a formal language. It was first described by the linguist Chomsky [1883] in 1956 [1883, 1884, 1885] and distinguishes four different classes of grammars. Starting with an unbounded grammar (type-0), more and more restrictions are imposed on the allowed production rules. Hence, each type contains all grammar types on higher levels fully.

Grammar	Allowed Rules	Languages
Type-0	$\alpha \rightarrow \beta, \alpha, \beta \in V^*, \alpha \neq \varepsilon$	recursive enumerable
Type-1	$\alpha A \beta \rightarrow \alpha \gamma \beta, A \in N, \alpha, \beta, \gamma \in V^*, \gamma \neq \varepsilon$	context-sensitive (CSG)
Type-2	$A \rightarrow \gamma, A \in N, \gamma \in V^*$	context-free ()
Type-3	$A \rightarrow aB$ (right-regular) or $A \rightarrow Ba$ (left-regular), $A \rightarrow a, A, B \in N, a \in \Sigma$	regular

Table 29.2: The Chomsky Hierarchy

Table 29.2 illustrates the Chomsky hierarchy. As already mentioned, V is the set containing all terminal and non-terminal symbols and V^* is its Kleene closure.

29.3.3 Derivation Trees

A derivation tree⁸² is a common way to describe how a sentence in a context-free language can be derived from the start symbol of a given generative grammar. The inner nodes of a derivation tree are the non-terminal symbols in N , the root is the start symbol S , and the leaves are the terminal symbols from the alphabet Σ . Each edge constitutes one expansion according to a production of the grammar.

Assume an example grammar $G = (N, \Sigma, P, S)$ with $N = \{T\}$, $\Sigma = \{1, +, a\}$, $S = T$, and the productions P as defined the below.

```

1  T  →  T+T
2  T  →  1
3  T  →  a

```

Listing 29.2: An example context-free generative grammar G .

⁸¹ http://en.wikipedia.org/wiki/Formal_grammar [accessed 2007-07-03]

⁸² http://en.wikipedia.org/wiki/Context-free_grammar#Derivations_and_syntax_trees [accessed 2007-07-16]

With this grammar we can construct the following sentence:

```

1  T      →  T+T
2  T+T    →  T+T+T
3  T+T+T  →  a+T+T
4  a+T+T  →  a+1+T
5  a+1+T  →  a+1+a

```

Listing 29.3: An example expansion of G .




Fig. 29.12: The derivation of the example expansion of the grammar G .

Figure 29.12 illustrates the derivation tree that belongs to this example expansion of the example grammar G .

29.3.4 Backus-Naur Form

The Backus-Naur (BNF) form⁸³ is a metasyntax used to express context-free grammars [1886, 1887]. Such Chomsky Type-2 grammars are the theoretical basis of most common programming languages and data formats, like for example C and XML⁸⁴. The BNF allows specifying production rules in simple, human and machine-understandable manner.

In BNF specifications, each rule consists of two parts: a non-terminal symbol on the left-hand side and an expansion on the right-hand side. Non-terminal symbols are contained in arrow brackets and terminal symbols are written plain. For expansions, the BNF provides two constructs: a sequence of symbols and the alternative which is denoted with a pipe character “|”.

Beginning with the start symbol $S = S$, the example below allows us to generate arbitrary natural numbers from $\mathbb{N}$. A **nonZero** is either 1,2,..., or 9 and a normal **number** may also be zero. A natural number is either a **nonZero** number or a natural number with a **number** at the end. Notice that expanding **nonZero** will always lead to the first digit being a non-zero digit since a fully expanded rule cannot contain any variables (non-terminal symbols). As start symbol $S = S$, we use **natural**.

```

1  <nonZero> ::= 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
2  <number>  ::= 0 | <nonZero>
3  <natural>  ::= <nonZero> | <natural> <number>
4  <S>       ::= <natural>

```

Listing 29.4: Natural numbers – a small BNF example.

⁸³ http://en.wikipedia.org/wiki/Backus%E2%80%93Naur_form [accessed 2007-07-03]

⁸⁴ <http://www.w3.org/TR/2006/REC-xml-20060816/> [accessed 2007-07-03]

29.3.5 Extended Backus-Naur Form

The extended Backus-Naur form⁸⁵ is an extension of the BNF metasyntax that provides additional operators and simplifications [1888, 1889, 1880]. Unlike the Backus-Naur form, the terminal symbols are included in quotation marks and the non-terminal symbols are written without arrow brackets. The items of sequences *can* now be separated by commas and each rule ends with a semicolon. The EBNF adds options, which are denoted by square brackets. The sequence inside such options may either occur zero or one time in the expanded rule. Curly brackets define expressions that can be left away or repeated arbitrary often during expansion.

The example below demonstrates the application of these new features by providing a grammar for natural numbers equal to the one shown for the BNF. The rules `natural` and `natural2` are equivalent. Here we also specify a rule for all integer numbers from $\mathbb{Z}$ by prefixing a natural number with an optional `-`.

```

1 nonZero    ::= "1" | "2" | "3" | "4" | "5" | "6" | "7" |
2              "8" | "9" ;
3 number     ::= "0" | nonZero ;
4 natural    ::= nonZero | natural, number ;
5 natural2   ::= nonZero | nonZero, {number} ;
6 integer    ::= ["-"], natural | "0" ;
7 S          ::= integer ;

```

Listing 29.5: Integer numbers – a small EBNF example.

The ISO norm ISO/IEC 14977 [1888] for EBNF defines additional extension mechanisms which we will not discuss here.

29.3.6 Attribute Grammars

An attribute grammar⁸⁶ (AG) is a context-free grammar enriched with attributes, rules, and conditions for these attributes [1890, 1891, 1892, 1893]. With attributes attached to non-terminal symbols, it becomes possible to provide context-sensitive information. Attribute grammars are often used in compilers to check rules that cannot be validated with the means of mere context-free grammars. With attribute grammars, syntax trees can be translated directly into intermediate languages or into code for some specific machine.

An attribute grammar $AG = (G, A, R)$ consists of three components:

1. a context-free grammar G , where $G = (N, \Sigma, P, S)$ as specified in Definition 29.47 on page 503,
2. a finite set of attributes A where each attribute $a \in A$ has a set of possible values $a = \{a_1, a_2, \dots, a_n\}$, and
3. a set of semantic rules R .

To each grammar symbol $X \in V$, a finite set of attributes $A(X) \subseteq A$ is associated. This set is partitioned into two disjoint subsets, the inherited attributes $I(X) \subseteq A(X)$ and the synthesized attributes $T(X) \subseteq A(X)$. The value of a synthesized attribute is determined by the attributes attached to the children of the symbol in the derivation tree it is assigned to. Inherited attributes get their value from the parent or siblings of the symbols they belong to. In the original definition by Knuth [1890], this was the other way round but the form discussed here has prevailed [1893]. The start symbol $S \in N$ and the terminal symbols in Σ do not have inherited attributes ($I(S) = \emptyset$, $\forall \sigma \in \Sigma \Rightarrow I(\sigma) = \emptyset$).

A good example for synthesized attributes is given in [1894] from where I will borrow. AGs are most often not used as generative grammars but as guidelines for parsers that read for instance source code of a programming language.

⁸⁵ <http://en.wikipedia.org/wiki/Ebnf> [accessed 2007-07-03]

⁸⁶ http://en.wikipedia.org/wiki/Attribute_grammar [accessed 2007-07-03]

Let us consider a simple grammar for integer mathematics with the two expressions + and *.

```

1 E ::= F      "+" E |
2   F
3 F ::= integer "*" F |
4   integer

```

Listing 29.6: A simple context-free grammar.

For each symbol X in V let $X.val$ be the numeric value associated with it. For terminal symbols of the type `integer`, this is simply the lexeme provided by the lexical analyzer. The two other terminal characters + and * have no value assigned. The values of the non-terminal symbols E and F should be the results of the expressions defined by them. These attributes are computed (synthesized) by the semantic rules from the attributes of their child nodes.

Production	Rule
$E ::= F \quad "+" \quad E \mid$	$E.val = F.val \quad + \quad E_2.val$
F	$E.val = F.val$
$F ::= integer \quad "*" \quad F \mid$	$F.val = integer.val \quad * \quad F_2.val$
$integer$	$F.val = integer.val$

Listing 29.7: A small example for attribute grammars.




Fig. 29.13: An instantiation of the grammar from Listing 29.7.

Figure 29.13 illustrates the derivation tree of a sentence of the simple attribute grammar from Listing 29.7. The non-terminal symbols are sometimes annotated with subscript numbers (like E_2) which have no meaning and only serve for clarity. While this Listing 29.7 is an example for the usage of synthesized attributes, symbol tables used in compilers are instances of inherited attributes.

A special form of attribute grammars, the reflective attribute grammar, is the basis of the Gads 2 Genetic Programming system discussed in Section 4.5.7 on page 155.

L-Attributed Grammars

L-attributed grammars⁸⁷ are a class of attribute grammars that can be parsed in one left-to-right traversal of the abstract syntax tree (see Section 4.1.1 on page 128). Such grammars are the foundations for many programming languages and allow convenient top-down parsing⁸⁸.

⁸⁷ http://en.wikipedia.org/wiki/L-attributed_grammar [accessed 2007-07-04]

⁸⁸ http://en.wikipedia.org/wiki/Top-down_parsing [accessed 2007-07-04]

S-Attributed Grammars

An attribute grammar is called S-attributed⁸⁹ if it allows only synthesized attributes [1895]. Because of this restriction, such grammars can be parsed top-down as well as directly bottom-up⁹⁰ and are supported by various tools like Bison⁹¹ and Flex⁹².

29.3.7 Extended Attribute Grammars

Extended Attribute Grammars developed by Watt [1896] and Madsen [1897] (EAGs) are a form of attribute grammars where the semantic (attribute-concerning) rules are no longer separated from the syntax productions [1010]. Instead, both are combined in a declarative form where each non-terminal symbol is accompanied by its attributes listed in a predetermined order. The new syntax for non-terminal symbols is

```
1 <n ↑a ↑b ↑c ...>
```

Listing 29.8: Syntax of an Extended Attribute Grammar symbol.

While $n \in N$ is a non-terminal symbol and a , b , and c are values of attributes α , β , and γ defined as expressions over their respective attribute value domain. In an extended attribute grammar, we can define a set of inherited attributes $I(n)$ and a set of synthesized attributes $T(n)$ for each non-terminal symbol n . In our example Listing 29.8, $\uparrow$ therefore has to be replaced with either $\downarrow$ which means that the following attribute is inherited ($\downarrow a \Leftrightarrow \alpha \in I(n)$) or $\uparrow$ denoting a synthesized attribute ($\uparrow a \Leftrightarrow \alpha \in T(n.\text{parent})$) where $n.\text{parent}$ stands for the parent node of n in the derivation tree. Terminal symbols cannot have attributes. Again, notice that the identifiers a , b , and c do not denote the attribute names but expressions that define their values. Attributes in EAGs are solely identified by their position in the non-terminal symbol specifications.

How this approach works is best understood when again, using a simple example borrowed from [1010]. Assume the grammar $G_1 = (N, \Sigma, P, S)$ with the non-terminal symbols $N = \{S, X, Y, Z\}$, the alphabet $\Sigma = \{x, y, z, \varepsilon\}$, productions P as defined below and the start symbol $S = S$. Additionally, X , Y , and Z are equipped with one synthesized attribute $v \in \mathbb{N}_0$.

```
1 <S>      ::= <X ↑v><Y ↑v><Z ↑v>
2 <X ↑v+1> ::= <X ↑v>"x"
3 <Y ↑v+1> ::= <Y ↑v>"y"
4 <Z ↑v+1> ::= <Z ↑v>"z"
5 <X ↑0>    ::= ε
6 <Y ↑0>    ::= ε
7 <Z ↑0>    ::= ε
```

Listing 29.9: The small example G_1 for Extended Attribute Grammars.

In the listing, below a typical expansion of G_1 is illustrated. Since the same attribute v is attached to all three non-terminals X , Y , and Z , the terminal symbols x , y , and z will always occur equally often. The context-sensitive grammar specified in Listing 29.9 thus defines sentences in the form $x^n y^n z^n$.

```
1 <S>  → <X ↑2><Y ↑2><Z ↑2> → <X ↑1>x<Y ↑2><Z ↑2>
2      → <X ↑0>xx<Y ↑2><Z ↑2> → xx<Y ↑2><Z ↑2>
3      → xx<Y ↑1>y<Z ↑2> → xx<Y ↑0>yy<Z ↑2>
4      → xxyy<Z ↑2> → xxyy<Z ↑1>z → xxyy<Z ↑0>zz
5      → xxyyzz
```

Listing 29.10: A typical expansion of G_1 .

⁸⁹ http://en.wikipedia.org/wiki/S-attributed_grammar [accessed 2007-07-04]

⁹⁰ http://en.wikipedia.org/wiki/Bottom-up_parsing [accessed 2007-07-04]

⁹¹ http://en.wikipedia.org/wiki/GNU_Bison [accessed 2007-07-04]

⁹² http://en.wikipedia.org/wiki/Flex_lexical_analyser [accessed 2007-07-04]

Another example for extended attribute grammars, again borrowed from [1010], demonstrates the specification of binary numbers. We can define a grammar $G_2 = (N, \Sigma, P, S)$ for all binary numbers. In this grammar, the start symbol S will have an attribute including the value of number represented by the generated sentence. Here we need three non-terminal symbols $N = \{S, T, B\}$ and only two terminal symbols $\Sigma = \{0, 1\}$. The productions P are specified as follows:

```

1 <S ↑b>      ::= <T ↓0 ↑b>
2 <T ↓a ↑b>    ::= <B ↓a ↑b>
3 <T ↓a ↑b+c>  ::= <T ↓a+1 ↑b><B ↓a ↑c>
4 <B ↓a ↑0>    ::= "0"
5 <B ↓a ↑2^a>  ::= "1"

```

Listing 29.11: An extended attribute grammar G_2 for binary numbers.

Figure 29.14 illustrates one possible expansion of the start symbol $S = S$ with the extended attribute grammar G_2 . As you can see, S has attached the (decimal) value 10 corresponding to the (binary) value 1010 of the binary string represented by the generated sentence.




Fig. 29.14: One possible expansion of the example grammar G_2 .

Extended Attribute Grammars are sufficient to specify the syntax and semantics of many programming languages [1896].

29.3.8 Adaptive Grammars

Definition 29.48 (Adaptive Grammar). An Adaptive Grammar⁹³ $G = (N, \Sigma, P, S)$ is a formal grammar developed by Shutt [1010] in which the set of non-terminal symbols N , the set of terminal symbols Σ and the set of productions P may vary during parsing.

⁹³ http://en.wikipedia.org/wiki/Adaptive_grammar [accessed 2007-07-13]

Shutt [1010] furthermore discusses Recursive Adaptive Grammars (RAG) which are a Turing complete formalism but yet retain the elegance of context-free grammars.

29.3.9 Christiansen Grammars

Christiansen [1013] introduces an adaptable grammar model that combines Extended Attribute Grammars with the ability to adapt according to Definition 29.48 [1013, 1898, 1899].

Unfortunately, Christiansen [1900] calls his adaptable attribute grammars “generative grammars” [1900, 1901] which has already another meaning (see Section 29.3.2 on page 502). We therefore resort to the term “Christiansen Grammars” coined by Shutt [1010] from whom we again will borrow the examples. As described in [1013, 1899], a Christiansen Grammar is an Extended Attribute Grammar where the first attribute of each non-terminal symbol $n \in N$ is an inherited Christiansen Grammar itself. This attribute is called *language attribute* and the expansion of the non-terminal symbol it belongs to must be done according to the grammar represented by it.

```
1 <n ↓g ↑a ↑b ...>
```

The statement $X\langle n \downarrow g \uparrow \dots \rangle Z ::= XYZ$ (with $X, Y, Z \in V$ and $n \in N$) hence only holds if $\langle n \downarrow g \uparrow a \dots \rangle ::= Y$ according to the grammar attribute g .

Let us start with a simple example grammar $G_3 = (N, \Sigma, P, S)$ with the non-terminal symbols **alpha-list** and **alpha**, the Latin alphabet as set of terminal symbols Σ , the **alpha-list** as start symbol S and the set of productions P as specified below.

```
1 <alpha-list ↓g ↑w>      ::= <alpha ↓g ↑w>
2 <alpha-list ↓g ↑w1ow2> ::= <alpha ↓g ↑w1><alpha-list ↓g ↑w2>
3 <alpha ↓g ↑"a">        ::= "a"
4 ...
5 <alpha ↓g ↑"z">        ::= "z"
```

Listing 29.12: A Christiansen Grammar creating character strings.

It clearly generates the character strings over the Latin alphabet. The start symbol has two attributes: The inherited Christiansen Grammar g will be handed down to all generated symbols. The attribute w on the other hand is synthesized from these symbols and contains the character string generated.

Basing on this grammar which still is a mere EAG in principle, we build the Christiansen Grammar $G_4 = (N, \Sigma, P, S)$ for a subset of the C (or Java) programming language where all value assignments are valid:

```
1 ...
2 <program ↓g0>          ::= "{"<decl-list ↓g0 ↑g1>
3                          <stmt-list ↓g1>"}"
4 <decl-list ↓g ↑g>      ::= ε
5 <decl-list ↓g0 ↑g2>    ::= <decl ↓g0 ↑g2><decl-list ↓g1 ↑g2>
6 <decl ↓g ↑g+new-rule> ::= "int" <alpha-list ↓g ↑w> ";"
7       where new-rule is <id ↓h> ::= w
8 <stmt-list ↓g>        ::= ε
9 <stmt-list ↓g>        ::= <stmt ↓g><stmt-list ↓g>
10 <stmt ↓g>            ::= <id ↓g> "=" <id ↓g> ";"
```

Listing 29.13: Christiansen Grammar for a simple programming language.

Whenever the non-terminal symbol **decl** is expanded, it also adds a new rule to the grammar. By introducing a new production for the symbol **id**, the declared variable becomes available in **stmt** since the grammar is synthesized upwards to the production for **program** and then inherited downwards into **stmt-list**. A more thorough example of Christiansen Grammar in the context of Genetic Programming can be found in Listing 4.7.

29.3.10 Tree-Adjoining Grammars

Tree-adjoining grammars⁹⁴ (TAG, also called tree-adjunct grammars) are another method for defining formal grammars which has been developed by Joshi [1902]. [1903, 1904] Different from BNF and EBNF, they are based on trees instead of plain strings. The inner nodes of the (fully expanded) trees correspond to non-terminal symbols and the leaves to terminal symbols.




Fig. 29.15: An example TAG tree.

The simple TAG tree illustrated in Figure 29.15 is borrowed from [1904] as well as some of the following examples. The tree structure of tree-adjoining grammars has one striking advantage compared to the flat rules in context-free grammars: the increased *domain of locality* [1903]. If we process for example an EBNF rule, we can only expand the non-terminal symbols at our current “level” of the derivation tree. Below we show that a text in an EBNF grammar similar to the one of Figure 29.15 could be resolved step by step. The variable `VP` expanded in line 8 for instance cannot be accessed or modified in line 10 anymore, although it is clearly part of the sentence construction.

```

1  S  ::= NP, VP ;
2  NP ::= "John" | "Lyn" ;
3  VP ::= V, NP ;
4  V  ::= "likes" ;
5
6  text → S
7  text → NP VP
8  text → "John" VP
9  text → "John" V NP
10 text → "John" "likes" NP
11 text → "John" "likes" "Lyn"

```

Listing 29.14: Another simple context-free grammar.

The extended domain of locality (EDL) in TAG trees is utilized with the two modification operators *substitution* and *adjunction*.

We can substitute a tree β into a tree α if there is a non-terminal leaf symbol ν in α that has the same label as the root of β . The stump of β then replaces the node ν in α . In Figure 29.16 we outline how two trees β_1 and β_2 are substituted into a TAG tree α and a new tree α' is created.

Substitution is equivalent to the non-terminal expansion in BNF. The adjunction operator however adds access to the aforementioned layers which are buried in context-free grammars. In order to perform an adjunction, the tree α has to include one non-terminal symbol ν at some random place. The root of the *auxiliary tree* is also labeled with ν and

⁹⁴ http://en.wikipedia.org/wiki/Tree-adjoining_grammar [accessed 2007-07-03]




Fig. 29.16: An example for the substitution operation.

so is at least one of its leafs. We now can replace the node marked with ν in α with tree β . Whatever was attached to ν before now replaces the leaf node ν in β . The leaf node ν in β often is additionally marked with an asterisk (*). Figure 29.17 sketches such a replacement, with the result that the new sentence α' now contains the word "really".

With adjunction, TAGs are somewhere in between context-sensitive and context-free grammars. In the definition of a tree-adjoining grammar $G = (N, \Sigma, A, I, S)$, A is the set of auxiliary trees to be used in the adjunction operations. I is the set of initial trees that can be substituted into existing trees. The union of I and A , $E = I \cup A$ is called the set of elementary trees and replaces the set of productions P used in Chomsky grammars. N and Σ retain their meaning as set of non-terminal and terminal symbols respectively. Trees with the non-terminal symbol $X \in N$ as root are called X -type trees. $S \in N$ denotes the starting symbol and there must be at least one S -type elementary tree.

Definition 29.49 (Lexicalized Tree-Adjoining Grammars). A lexicalized tree-adjoining grammar (LTAG) is a tree-adjoining grammar where each elementary tree $t \in E$ contains a terminal symbol $X \in \Sigma$. Although they are more restricted, LTAGs are equivalent to TAGs.

A discussion on derivation trees of tree-adjoining grammars can be found in Section 4.5.9 on page 159.




Fig. 29.17: An example for the adjunction operation.

29.3.11 S-expressions

S-expressions⁹⁵ (where S stands for symbolic) or *sexp* are data structures for presenting complex data. They are probably best known for their usage in the Lisp⁹⁶ [1905, 1906, 1907] and Scheme⁹⁷ [1908] programming languages. Their most common feature is that they are parenthesized prefix notations (often also known as Polish notation⁹⁸).

In 1997, Rivest [1909] handed in a standardization draft for S-expressions to be considered for publication as RFC. It was, however, never approved but still is the foundation for many other publications and RFCs.

```

1 (defun fibonacci (N)
2   (if (or (zerop N) (= N 1))
3       1
4       (+ (fibonacci (- N 1)) (fibonacci (- N 2)))))

```

Listing 29.15: A small Lisp-example: How to compute Fibonacci numbers.

⁹⁵ <http://en.wikipedia.org/wiki/S-expression> [accessed 2007-07-03]

⁹⁶ http://en.wikipedia.org/wiki/Lisp_programming_language [accessed 2007-07-03]

⁹⁷ http://en.wikipedia.org/wiki/Scheme_%28programming_language%29 [accessed 2007-07-03]

⁹⁸ http://en.wikipedia.org/wiki/Polish_notation [accessed 2007-07-04]

Part V

Appendices

A

Abbreviations and Symbols

This section contains a list of the abbreviations any symbols used in the context of this book. As you will see, some symbols have multiple meanings and you have to guess the right one using the context – otherwise we would run out of letters.

TODO

Symbols

$\vdash$	denotes the Pareto dominance relation, $x_1 \vdash x_2$ means that x_1 is better than x_2 in the Pareto sense (see Definition 1.13 on page 13)
$\succ$	denotes the prevalence relation, $x_1 \succ x_2$ means that x_1 is better than x_2 in a user-defined sense (see Definition 1.17 on page 20)
$\text{len}(l)$	the number of elements in the list, set, or tuple l
$n!$	the factorial of the natural number n
$\bar{a}$	the arithmetic mean of all elements a (see Definition 27.29 on page 424)
$\binom{M}{n}$	the number of combinations of n out of M elements (see Section 27.1.1 on page 418)
$(M)_n$	the number of permutations of n out of M elements (see Section 27.1.1 on page 419)
$\ x\ _n$	the n -norm of x (see Definition 28.9 on page 470)
$(1, 2)^T$	the transposition of a vector, i. e., $(1, 2)^T = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$
$\alpha \circ \beta$	the string β attached (concatenated) to the string α (see Definition 29.44 on page 502)
$a, \mathbf{a}$	Notation: a is a scalar value whereas $\mathbf{a}$ is a vector
∇	the gradient of a given function (see Definition 1.41 on page 34)

α

α a character from a given alphabet Σ (see Definition 29.42 on page 502)

γ, Γ

γ	a confidence coefficient/probability (see Section 27.6.3 on page 454)
γ	the Euler-Mascheroni constant (see Equation 27.304 on page 466)
γ_1	the skewness (see Definition 27.40 on page 427)
γ_2	the kurtosis (see Definition 27.41 on page 427)
$\gamma(a, b)$	the lower incomplete gamma function of a and b
$\Gamma(x)$	the gamma function of x (see Section 27.8.1 on page 466)
$\Gamma(a, b)$	the upper incomplete gamma function of a and b

ϵ

ϵ	the empty string (see Definition 29.43 on page 502)
$\epsilon(\tilde{\theta})$	the absolute, signed error of $\tilde{\theta}$ (see Definition 27.52 on page 450)
$\epsilon(f)$	multiple absolute errors, grouped to a vector, of f (see Equation 27.246 on page 453)

θ, Θ

θ	some stochastic parameter of a random distribution or experiment (see Definition 27.50 on page 450)
$\tilde{\theta}$	the estimate of the stochastic parameter θ computed using a data sample (see Definition 27.50 on page 450)
$\Theta(f)$	the Θ notation; the set $\Theta(f)$ of the function f (see Definition 29.15 on page 483)

μ

μ	the expected value of a random variable X (see Definition 27.27 on page 424)
μ'_k	the k^{th} statistical moment (see Definition 27.37 on page 426)
$\mu'_k(c)$	the k^{th} moment about a value c (see Definition 27.36 on page 426)
μ_k	the k^{th} (central) moment about the mean (see Definition 27.38 on page 426)
$\mu_{\sigma, k}$	the k^{th} standardized moment (see Definition 27.39 on page 426)

ρ

$\rho(x)$	a density estimation function of x (see Section 27.6.4 on page 457)

$\rho_{cr}(a, A)$	the crowding distance density estimate (see Section 27.6.4 on page 457)
$\rho_{nn,k}(a, A)$	the k^{th} nearest neighbor density estimate (see Section 27.6.4 on page 457)
ρ	the autocorrelation function (see Definition 1.53 on page 43)

 π, Π

π_λ	stands for the Poisson distribution (see Section 27.3.2 on page 431)

 σ, Σ

σ	the standard deviation of a probability distribution (see Definition 27.33 on page 425)
σ^2	the variance of a probability distribution (see Definition 27.30 on page 424)
Σ	an alphabet, for example of a given grammar; (see Definition 29.42 on page 502)
Σ^l	all words of a given length l over an alphabet Σ (see Section 29.3.1 on page 502)
Σ^*	all words over an alphabet Σ (see Section 29.3.1 on page 502)

 τ

τ	an index identifying the τ^{th} evaluation of an individual (see Definition 1.42 on page 34)

 ϕ, Φ

φ_X	the characteristic function of a probability distribution of a random variable X
$\Phi(x)$	the cdf of the standard normal distribution (see Definition 27.48 on page 437)
$\Phi^{-1}(b)$	the inverse of the cdf of the standard normal distribution (see Definition 27.48 on page 437)
$\Phi(x, \tau)$	the problem landscape; the probability of reaching element $x \in \mathbb{X}$ until the τ^{th} individual evaluation (see Definition 1.38 on page 29)

 ψ, Ψ

$\psi(a)$	the digamma function of a

 ω, Ω

$\omega(f)$	the small- ω notation; the set $\omega(f)$ of the function f (see Definition 29.17 on page 484)
ω	an elementary event – an element of the sample space Ω of a random experiment (see Definition 27.2 on page 417)
$\Omega(f)$	the big- Ω notation; the set $\Omega(f)$ of the function f (see Definition 29.14 on page 483)
Ω	the sample space of a random experiment (see Definition 27.3 on page 417)

A

a	an element of the data sample A (see Section 27.2 on page 423); a is often also used to as name for ordinary set elements
a	an element of a set to which a clustering algorithm is to be applied (see Chapter 28 on page 467)
a	the probability that the real value of the estimated parameter lies not inside the confidence interval (see Section 27.6.3 on page 454)
A	a data sample obtained from random experiments (see Section 27.2 on page 423); A is often also used to as name for ordinary sets
A	a set to which a clustering algorithm is to be applied (see Chapter 28 on page 467)
A	the set of attributes of an attribute grammar (see Section 29.3.6 on page 505)
A	the set of auxiliary trees of a tree-adjoining grammar (see Section 29.3.10 on page 510)
$\text{addListItem}(l, q)$	insert the element q at the end of a list l (see Definition 26.16 on page 411)
adjacent	the adjacency relation of elements in the search or problem space (see Definition 1.29 on page 25 and Definition 1.31 on page 26)
$AG = (G, A, R)$	an attribute grammar AG consisting of a formal grammar G , a set of attributes A , and a set of rules R (see Section 29.3.6 on page 505)
$\text{appendList}(l_1, l_2)$	a shortcut for adding all the elements of a list l_2 to a list l_1 (see Definition 26.19 on page 411)

B

$\mathbb{B}$	the Boolean space $\mathbb{B} = \{\text{true}, \text{false}\}$
$B(a, b)$	the Beta function of a and b
$B(n, p)$	stands for the binomial distribution Section 27.3.3
$\text{Bias}(\tilde{\theta})$	the bias of $\tilde{\theta}$ (see Definition 27.53 on page 450)

C

$\mathbb{C}$	the complex numbers (see Section 26.3 on page 408)

$C(M, n)$	the number of combinations of n out of M elements (see Section 27.1.1 on page 418)
c_V	the coefficient of variation (see Definition 27.34 on page 425)
c	one cluster – part of the result of a clustering algorithm (see Chapter 28 on page 467)
C	the set of all clusters – the result of a clustering algorithm (see Chapter 28 on page 467)
$\mathfrak{C}$	all possible cluster partitions (see Chapter 28 on page 467)
$\mathfrak{C}_k$	all possible cluster partitions which consist of exactly k clusters (see Chapter 28 on page 467)
$\text{centroid}(c)$	the centroid of the cluster c (see Definition 28.4 on page 469)
$\text{cluster}(A)$	a clustering algorithm applied to the set A (see Definition 28.2 on page 468)
cmp	a comparator function (see Definition 1.15 on page 19)
cmp_F	a comparator function used by prevalence optimization (see Definition 1.16 on page 20)
$\text{countItemOccurrences}(x, l)$	the number of occurrences of the element x in the list l (see Definition 26.20 on page 411)
$\text{cov}(X, Y)$	the covariance of X and Y (see Definition 27.35 on page 426)
$\text{createList}(n, q)$	create a new list of the length n filled with the item q (see Definition 26.14 on page 411)

D

DX	the standard deviation of a random variable X (see Definition 27.33 on page 425)
D^2X	the variance of a random variable X (see Definition 27.30 on page 424)
$\text{deleteListItem}(l, i)$	create a new list by removing the element at index $0 \leq i < \text{len}(l)$ from the list l (see Definition 26.17 on page 411)
$\text{deleteListRange}(l, i, c)$	create a new list by removing c elements beginning at index $0 \leq i < \text{len}(l)$ from the list l (see Definition 26.18 on page 411)
$\text{dist}(a, b)$	a measure for the distance between a and b
$\text{dist}_{cr}^\rho(a, A)$	the crowding distance (see Section 27.6.4 on page 457)
$\text{dist}_{nn,k}^\rho(\text{dist}, a, A)$	the k^{th} nearest neighbor distance (see Definition 27.62 on page 457)
$\text{dist}_{eucl}(a, b)$	the Euclidian distance between two elements a and b (see Definition 28.8 on page 470)
$\text{dist}_{n,n}(a, b)$	the n -norm distance of two elements a and b (see Equation 28.17 on page 470)
$\text{dist}_{Ham}(a, b)$	the Hamming distance of two elements a and b (see Definition 28.6 on page 469)
$\text{dist}_{Man}(a, b)$	the Manhattan distance of two elements a and b (see Definition 28.7 on page 469)
$\text{dist}_{max}(c_1, c_2)$	the maximum distance (complete linkage) between two clusters c_1 and c_2 (see Section 28.1.4 on page 471)

$\text{dist}_{\min}(c_1, c_2)$	the minimum distance (single linkage) between two clusters c_1 and c_2 (see Section 28.1.4 on page 471)
$\text{dist}_{\text{avg}}(c_1, c_2)$	the average distance (average linkage) between two clusters c_1 and c_2 (see Section 28.1.4 on page 471)
$\text{dist}_{\text{var}}(c_1, c_2)$	a distance measure between two clusters c_1 and c_2 denoting the variance change if they were merged (see Section 28.1.4 on page 471)
$\text{dist}_{\text{cent}}(c_1, c_2)$	the distance between the centroids of two clusters c_1 and c_2 (see Section 28.1.4 on page 471)
$\text{dist}_{\text{nuc}}(c_1, c_2)$	the distance between the nuclei two clusters c_1 and c_2 (see Section 28.1.4 on page 471)

E

E	the set of elementary trees of a tree-adjoining grammar (see Section 29.3.10 on page 510)
$\text{error}_c(c)$	the clustering error of a single cluster c (see Equation 28.27 on page 472)
$\text{error}_p(C)$	the clustering error of a set of clusters C (see Equation 28.28 on page 472)
$\text{errorBurst}(m, \varepsilon)$	introduces an error burst into the message m with probability ε (see Algorithm 29.2 on page 500)
EX	the expected value of a random variable X (see Definition 27.27 on page 424)
$\exp(\lambda)$	stands for the exponential distribution (see Section 27.4.3 on page 440)

F

F	the set of objective functions f applied to a given problem (see Chapter 1 on page 3)
f	an objective function, i. e., $f \in F$ (see Definition 1.1 on page 3)
F_X	a cumulative distribution function of a random variable X (see Definition 27.18 on page 421)
f_X	a probability mass function of a <i>discrete</i> random variable X (see Section 27.1.7 on page 422)
f_X	a probability density function of a <i>continuous</i> random variable X (see Section 27.1.8 on page 422)
${}_2F_1(a, b, c, d)$	the hypergeometric function of a, b, c, d

G

$\mathbb{G}$	the genome, the search space – the elements on which an search- or optimization algorithms work (see Definition 1.21 on page 23)
g	a genotype, an element of the search space, i. e., $g \in \mathbb{G}$ (see Definition 1.22 on page 23)

G	a subset of the search space, i. e., $G \subseteq \mathbb{G}$
G_1	the unbiased estimate of the skewness of a data sample (see Section 27.2.5 on page 427)
G_2	the unbiased estimate of the kurtosis of a data sample (see Section 27.2.5 on page 427)
$G = (N, \Sigma, P, S)$	a formal grammar G consisting of the non-terminal symbols N , the alphabet Σ , the productions P , and the start symbol S (see Definition 29.47 on page 503)
$G = (N, \Sigma, A, I, S)$	a tree-adjoining grammar G consisting of the non-terminal symbols N , the terminal symbols Σ , the auxiliary trees A , the initial trees I , and the start symbol S (see Section 29.3.10 on page 510)
gpm	the genotype-phenotype mapping (see Definition 1.30 on page 25)
gpm^{-1}	the inverse function to the genotype-phenotype mapping, does not necessarily exist (see Equation 1.29 on page 26)
grad	the gradient of a given function (see Definition 1.41 on page 34)

H

$h(A, n)$	the relative frequency: the fraction of the number of occurrences of an event A during n repetitions of a random experiment (see Definition 27.10 on page 419)
$H(A, n)$	the absolute frequency: how often an event A occurred during n repetitions of a random experiment (see Definition 27.9 on page 419)
H	the (information) entropy of a data sample or probability distribution (see Definition 27.46 on page 429)
h	the (differential) entropy of a data sample or probability distribution (see Definition 27.47 on page 429)

I

i	the imaginary unit (see Section 26.3 on page 408)
I	the inherited attributes of an attribute grammar (see Section 29.3.6 on page 505)
I	the set of initial trees of a tree-adjoining grammar (see Section 29.3.10 on page 510)
$I_{1-n}(a, b)$	the regularized incomplete Beta function of n , a , and b
$\text{Im}(z)$	the imaginary part of a complex number z (see Section 26.3 on page 408)
$\text{insertListItem}(l, i, q)$	create a new list by inserting one element q in a list l at the index $0 \leq i \leq \text{len}(l)$ (see Definition 26.15 on page 411)

K

kMeansModify	the k means modify function used by the k means clustering algorithm (see Section 28.2.2 on page 472)

kMeansCluster the k means clustering algorithm (see Section 28.2.2 on page 472)

L

$\mathbb{L}$	a formal language (see Definition 29.46 on page 502)
leaderCluster	the leader clustering algorithm (see Algorithm 28.4 on page 476)
$L[\tilde{\theta}]$	the likelihood of $\tilde{\theta}$ (see Section 27.6.1 on page 451)
linkageCluster	the linkage clustering algorithm family (see Section 28.2.4 on page 473)
listToSet(B)	transform a list B to a set (see Equation 26.50 on page 413)

M

M	the moment generating function of a probability distribution
max	the maximum element (context-dependent)
med	the median of a data sample or a probability distribution (see Definition 27.42 on page 427)
min	the minimum element (context-dependent)
mode	the mode of a data sample or a probability distribution (see Definition 27.45 on page 428)
$\text{MSE}(\tilde{\theta})$	the mean square error of $\tilde{\theta}$ (see Definition 27.55 on page 450)
$\widetilde{\text{MSE}}(\tilde{\theta})$	the estimated mean square error of $\tilde{\theta}$ (see Equation 27.234 on page 451)

N

$\mathbb{N}$	the natural numbers, <i>not</i> including 0: $\{1, 2, 3, 4, \dots\}$ (see Section 26.3 on page 408)
$\mathbb{N}_0$	the natural numbers including 0: $\{0, 1, 2, 3, \dots\}$ (see Section 26.3 on page 408)
N	the non-terminal symbols of a given grammar (see Definition 29.47 on page 503)
n	stands usually for a node in a network
$\mathcal{N}$	stands usually for a network, a distributed system of computational devices
nNearestNeighborCluster	the n nearest neighbor clustering algorithm (see Section 28.2.3 on page 472)
$N(\mu, \sigma^2)$	stands for the normal distribution (see Section 27.4.2 on page 437)
nucleus(c)	the nucleus of the cluster c (see Equation 28.21 on page 471)

O

$\mathbf{o}(f)$	the small- $\mathbf{o}$ notation; the set $\mathbf{o}(f)$ of the function f (see Definition 29.16 on page 483)

$\mathbf{O}(f)$	the big- $\mathbf{O}$ notation; the set $\mathbf{O}(f)$ of the function f (see Section 29.1.3 on page 482)
Op	the set of search operations (see Definition 1.26 on page 24)
searchOp	a search operation (see Definition 1.26 on page 24)

P

p	an individual - a tuple of a genotype $p.g \in \mathbb{G}$ and a phenotype $p.x \in \mathbb{X}$ (see Definition 1.36 on page 28)
$\mathcal{P}(A)$	the power set of the set A (see Definition 26.9 on page 410)
P	the productions of a given grammar (see Definition 29.47 on page 503)
$P(\omega)$	the probability of an event ω (see Definition 27.13 on page 420)
$P_\gamma(a, b)$	the regularized gamma function of a and b
$Perm(M, n)$	the number of permutations of n out of M elements (see Section 27.1.1 on page 419)
Pop	a population of individuals (see Definition 1.37 on page 28)
probit(a)	the probit function of a (see Definition 27.49 on page 439)

Q

$\mathbb{Q}$	the rational numbers, e.g., $\frac{1}{2}, -\frac{23}{7}, 11, \dots$ (see Section 26.3 on page 408)
quantile	a quantile of a data sample or probability distribution (see Definition 27.43 on page 428)

R

$\mathbb{R}$	the real numbers, e.g., $\frac{1}{2}, -\frac{23}{7}, 11, \pi, e, \dots$ (see Section 26.3 on page 408)
$\mathbb{R}^+$	the positive real numbers: $\mathbb{R}^+ = \{\forall x : x \in \mathbb{R} \wedge x \geq 0\}$ (see Section 26.3 on page 408)
$\mathbb{R}^n$	a n -dimensional space of all vectors composed of n real numbers (see Section 26.3 on page 408)
R	the rules of an attribute grammar (see Section 29.3.6 on page 505)
$R(a_1, a_2)$	the elements a_1 and a_2 are related by the binary relation R (see Section 26.7 on page 413)
random	a random number generator function (see Definition 27.64 on page 461)
random _{bs} (μ, σ)	a random number generator creating random numbers according to a bell-shaped distribution with mean μ and standard deviation σ (see Section 27.7.3 on page 464)
random _l (random, a, b)	limits the results of the random number generator random to the interval $[a, b]$ (see Section 27.7.2 on page 462)
random _n (μ, σ^2)	a random number generator function creating normally distributed random numbers with mean μ and variance σ^2 (see Definition 27.66 on page 461)

$\text{random}_u(a, b)$	a random number generator function creating random numbers uniformly distributed in $[a, b]$ (see Definition 27.65 on page 461)
range	the statistical property range (see Definition 27.26 on page 423)
$\text{Re}(z)$	the real part of a complex number z (see Section 26.3 on page 408)
$\text{removeListItem}(l, q)$	remove one occurrence of the item q from list l (see Definition 26.25 on page 413)

S

s	the unbiased estimate of the standard deviation of a data sample (see Definition 27.33 on page 425)
S	the start symbol of a given grammar (see Definition 29.47 on page 503)
$\mathbb{S}$	the solution space (see Definition 1.20 on page 23)
s	an element of the solution space
S	a subset of the solution space
S	the Sterling function (see Equation 28.6 on page 468)
s^2	the unbiased estimate of the variance of a data sample (see Definition 27.32 on page 425)
$\text{searchItem}_{as}(s, S)$	search an element s in an ascending sorted list S (see Definition 26.23 on page 412)
$\text{searchItem}_{ds}(s, S)$	search an element s in a descending sorted list S (see Definition 26.23 on page 412)
$\text{searchItem}_u(u, U)$	search an element u in an unsorted list U (see Definition 26.23 on page 412)
$\text{setToList}(A)$	transform the set A to a list (see Equation 26.49 on page 413)
$\text{sortList}_a(U, s)$	sort the list U ; bring it in ascending order according to the sorting criterion s (see Definition 26.22 on page 412)
$\text{sortList}_d(U, s)$	sort the list U ; bring it in descending order according to the sorting criterion s (see Definition 26.22 on page 412)
$\text{subList}(l, i, c)$	extract c elements from the list l beginning at index i and return them as a new list (see Definition 26.21 on page 412)
$\text{sum}(A)$	the sum of all elements in the list or set A (see Definition 27.28 on page 424)
$\text{sumSqs}(A)$	the sum of the squares of all elements in the list or set A (see Definition 27.31 on page 425)

T

t	the iteration counter (see Definition 1.43 on page 34)
T	the synthesized attributes of an attribute grammar (see Section 29.3.6 on page 505)
$\text{terminationCriterion}()$	the termination criterion telling the optimizer when to stop (see Definition 1.44 on page 35)

V

V	the set of all symbols (terminal and non-terminal) $V = \Sigma \cup N$ of a given grammar (see Definition 29.47 on page 503)
$\mathbb{V}$	the fitness space – the space that contains the fitness values (see Definition 1.35 on page 28)
$v(x)$	the fitness value assigned to an element x of the problem space (see Definition 1.35 on page 28)
V	a subset of the fitness space, i. e., $V \subseteq \mathbb{V}$

X

$\mathbb{X}$	the phenome, the problem space – the space that contains the possible solutions/optima to a given optimum (see Chapter 1 on page 3)
x	a solution candidate, an element of the problem space, i. e., $x \in \mathbb{X}$ (see Chapter 1 on page 3)
X	a subset of the problem space, i. e., $X \subseteq \mathbb{X}$ (see Chapter 1 on page 3)
X	a random variable (see Definition 27.16 on page 421)
$\hat{x}$	a maximum element of a given set (most often of the <i>explored part</i> of the problem space $\mathbb{X}$) (see Section 1.2.1 on page 7)
$\hat{X}$	a maximum subset of a given set (most often of the <i>explored part</i> of the problem space $\mathbb{X}$) (see Section 1.2.1 on page 7)
$\hat{x}_l$	a local maximum element of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.6 on page 7)
$\hat{\mathbf{x}}$	a global maximum element of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.9 on page 8)
$\hat{\mathbf{X}}$	a global maximum subset of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.9 on page 8)
$\tilde{x}$	a minimum element of a given set (most often of the <i>explored part</i> of the problem space $\mathbb{X}$) (see Section 1.2.1 on page 7)
$\tilde{X}$	a minimum subset of a given set (most often of the <i>explored part</i> of the problem space $\mathbb{X}$) (see Section 1.2.1 on page 7)
$\tilde{x}_l$	a local minimum element of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.7 on page 7)
$\tilde{\mathbf{x}}$	a global minimum element of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.10 on page 8)
$\tilde{\mathbf{X}}$	a global minimum subset of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.10 on page 8)
x^*	a optimum element of a given set (most often of the <i>explored part</i> of the problem space $\mathbb{X}$) (see Section 1.2 on page 7)
X^*	a optimum subset of a given set (most often of the <i>explored part</i> of the problem space $\mathbb{X}$) (see Section 1.2 on page 7)
x_l^*	a local optimum element of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.8 on page 7)

$\mathbf{x}^*$	a global optimum element of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.11 on page 8 and Section 1.2.2 on page 9)
$\mathbf{X}^*$	a global optimum subset of a given set (most often of the problem space $\mathbb{X}$) (see Definition 1.11 on page 8 and Section 1.2.2 on page 9)
$(\mathbb{X}, F, \mathbb{G}, Op, \text{gpm})$	an optimization problem (see Definition 1.34 on page 27)

Y

$\mathbb{Y}$	the objective space – the space that contains all possible objective value (see Definition 1.33 on page 27)
y	an element of the objective space, i. e., $y = F(x \in \mathbb{X}) \Leftrightarrow y \in \mathbb{Y}$
Y	a subset of the objective space, i. e., $Y \subseteq \mathbb{Y}$

Z

$\mathbb{Z}$	the integer numbers: $\{.., -2, -1, 0, 1, 2, ..\}$ (see Section 26.3 on page 408)

B

GNU Free Documentation License (FDL)

Version 1.2, November 2002

Copyright (C) 2000,2001,2002 Free Software Foundation, Inc.

51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

B.1 Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

B.2 Applicability and Definitions

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document

to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

B.3 Verbatim Copying

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this

License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in Section B.4.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

B.4 Copying in Quantity

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

B.5 Modifications

You may copy and distribute a Modified Version of the Document under the conditions of Section B.3 and Section B.4.3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

B.6 Combining Documents

You may combine the Document with other documents released under this License, under the terms defined in Section B.5 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

B.7 Collections of Documents

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

B.8 Aggregation with Independent Works

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of Section B.4 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

B.9 Translation

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of Section B.5. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version

of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (Section B.5) to Preserve its Title (Section B.2) will typically require changing the actual title.

B.10 Termination

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

B.11 Future Revisions of this License

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document

under the terms of the GNU Free Documentation License, Version 1.2

or any later version published by the Free Software Foundation;

with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.

A copy of the license is included in the section entitled "GNU

Free Documentation License".

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

C

GNU Lesser General Public License (LPGL)

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1]

C.1 Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages—typically libraries—of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

C.2 Terms and Conditions for Copying, Distribution and Modification

1. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightfor-

wardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

2. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

3. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

4. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that

refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

5. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

6. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

7. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library",

as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

- 8. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:
 - a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
 - b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.
- 9. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 10. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.
- 11. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with

or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

12. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

13. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
14. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.
15. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

C.3 No Warranty

1. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED,

INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

2. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

C.4 How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the library's name and an idea of what it does. Copyright (C) year name of author

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library 'Frob' (a library for tweaking knobs) written by James Random Hacker.

signature of Ty Coon, 1 April 1990 Ty Coon, President of Vice

That's all there is to it!

D

Credits and Contributors

In this section I want to give credits to whom they deserve. So its props to:

Stefan Achler

For working together at the 2007 DATA-MINING-CUP Contest.
See Section 21.1.2 on page 318 and [463].

Steffen Bleul

For working together at the 2006, 2007, and 2008 Web Service Challenge.
See Section 21.2 on page 327 and [1604, 1606, 1605, 1607]

Distributed Systems Group, University of Kassel

To the whole Distributed Systems Group at the University of Kassel for being supportive and coming over again and again with new ideas and helpful advices.
Each and every research project described in this book.

Gan Min

For careful reading and pointing out inconsistencies and spelling mistakes.
See especially the Pareto optimization area Section 1.2.4 on page 13.

Kurt Geihs

For being supportive and contributing to many of my research projects.
See for instance Section 21.1.2 on page 318 and [861, 862, 863, 1604, 1605]

Martin Göb

For working together at the 2007 DATA-MINING-CUP Contest.
See Section 21.1.2 on page 318 and [463].

Ibrahim Ibrahim

For careful proofreading.
See Chapter 1 on page 3.

Stefan Niemczyk

For careful proofreading and scientific support.
See for instance Chapter 1 on page 3, Section 20.2.6 on page 290, and [1546, 1547].

Roland Reichle

For fruitful discussions and scientific support.
See for instance Section 27.6.1 on page 451, Section 20.2.6 on page 290, and [1546].

Richard Seabrook

For pointing out mistakes in some of the set theory definitions.
See Chapter 26 on page 407

Hendrik Skubch

For fruitful discussions and scientific support.

See for instance Section 1.3 on page 21, Section 20.2.6 on page 290, and [1546].

Christian Voigtmann

For working together at the 2007 and 2008 DATA-MINING-CUP Contest.

See Section 21.1.2 on page 318 and [463].

Wibul Wongpoowarak

For valuable suggestions in different areas, like including random optimization.

See for instance Section 27.7.1 on page 460 [1293], Chapter 9 on page 213.

Michael Zapf

For helping me to make many sections more comprehensible.

See for instance Section 23.2.2 on page 367, Section 21.1.2 on page 318, Section 4.5 on page 147, and [1910, 1911].

E

Citation Suggestion

```
@book{W2008GOEB,  
  author      = {Thomas Weise},  
  title       = {Global Optimization Algorithms - Theory and Application},  
  year        = {2008},  
  month       = aug # {~12},  
  howpublished = {Online as e-book},  
  edition     = {Second},  
  institution = {University of Kassel, Distributed Systems Group},  
  organization = {University of Kassel, Distributed Systems Group},  
  publisher   = {Thomas Weise},  
  copyright   = {Copyright (c) 2006-2008 Thomas Weise, licensed under GNU FDL},  
  keywords    = {Global Optimization, Evolutionary Computation, Evolutionary Algorithms,  
                 Genetic Algorithms, Genetic Programming, Evolution Strategy, Learning  
                 Classifier Systems, Sigoa, Dgpf, Java},  
  
  language    = {en},  
  url         = {http://www.it-weise.de/},  
  url         = {http://www.sigoa.org/},  
  url         = {http://www.vs.uni-kassel.de/staff/researchers/weise/},  
  note        = {Online available at http://www.it-weise.de/}  
}
```

```
%0   Book  
%A   Weise, Thomas  
%T   Global Optimization Algorithms - Theory and Application  
%F   Thomas Weise: "Global Optimization Algorithms - Theory and Application"  
%I   Thomas Weise  
%D   Second  
%8   Second  
%7   Second  
%9   Online available as e-book at http://www.it-weise.de/
```

References

- [1] Jörg Heitkötter and David Beasley, editors. *Hitch-Hiker's Guide to Evolutionary Computation: A List of Frequently Asked Questions (FAQ)*. ENCORE (The Evolutionary Computation REpository Network), 1998. USENET: comp.ai.genetic. Online available at <http://www.cse.dmu.ac.uk/~rij/gafaq/top.htm> and <http://alife.santafe.edu/~joke/encore/WWW/> [accessed 2007-07-03].
- [2] Wikipedia. Wikipedia – The Free Encyclopedia, 2008. Online available at <http://en.wikipedia.org/> [accessed 2008-08-12].
- [3] R.A. The Rugged Man. Midnight Thud. In *Die, Rugged Man, Die*. Nature Sounds, November 16, 2004. ASIN: B00065GIKA.
- [4] Linus Pauling. *The Nature of the Chemical Bond*. Cornell Univ. Press, Ithaca, New York, 1960. ISBN 0-8014-0333-2. Online available at <http://osulibrary.oregonstate.edu/specialcollections/coll/pauling/> [accessed 2007-07-12].
- [5] Herbert Spencer. *The Principles of Biology*, volume 1. London & Edinburgh: Williams and Norgate, first edition, 1864 and 1867. Online available at <http://www.archive.org/details/ThePrinciplesOfBiology> [accessed 2007-08-05].
- [6] Charles Darwin. *On the Origin of Species*. John Murray, sixth edition, November 1859. Online available at <http://www.gutenberg.org/etext/1228> [accessed 2007-08-05].
- [7] Arnold Neumaier from the Computational Mathematics group, University of Vienna, Austria. Global optimization and constraint satisfaction. In I. Bomze, I. Emiris, Arnold Neumaier, and L. Wolsey, editors, *Proceedings of GICOLAG workshop (of the research project Global Optimization, Integrating Convexity, Optimization, Logic Programming and Computational Algebraic Geometry)*, December 2006. Online available at <http://www.mat.univie.ac.at/~neum/glopt.html> [accessed 2007-07-12].
- [8] Panos M. Pardalos, Nguyen Van Thoai, and Reiner Horst. *Introduction to Global Optimization*. Nonconvex Optimization and Its Applications. Springer, second edition, December 31, 2000. ISBN 978-0792367567. First edition: June 30, 1995, ISBN: 978-0792335566.
- [9] Zbigniew Michalewicz and David B. Fogel. *How to Solve It: Modern Heuristics*. Springer, second, revised and extended edition, December 2004. ISBN 978-3540224945.
- [10] Victor J. Rayward-Smith, Ibrahim H. Osman, Colin R. Reeves, and George D. Smith, editors. *Modern Heuristic Search Methods*. Wiley, December 1996. ISBN 978-0471962809.
- [11] Judea Pearl. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. The Addison-Wesley series in artificial intelligence. Addison-Wesley Pub (Sd), April 1984. ISBN 978-0201055948.
- [12] Kalyanmoy Deb. Evolutionary Algorithms for Multi-Criterion Optimization in Engineering Design. In *Evolutionary Algorithms in Engineering and Computer Science*, pages 135–161. John Wiley & Sons, Ltd, 1999. In collection [456]. Online avail-

- able at <http://citeseer.ist.psu.edu/deb99evolutionary.html> and <http://www.lania.mx/~ccoello/EM00/deb99.ps.gz> [accessed 2007-08-25].
- [13] Kalyanmoy Deb. *Multi-Objective Optimization Using Evolutionary Algorithms*. Wiley Interscience Series in Systems and Optimization. John Wiley & Sons, Inc., New York, NY, USA, May 2001. ISBN 978-0471873396.
 - [14] Carlos Artemio Coello Coello. An Updated Survey of Evolutionary Multiobjective Optimization Techniques: State of the Art and Future Trends. In *1999 Congress on Evolutionary Computation*, pages 3–13, 1999. In proceedings [502]. Online available at <http://citeseer.ist.psu.edu/coellocoello99updated.html> [accessed 2007-08-25].
 - [15] Carlos M. Fonseca and Peter J. Fleming. Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms – Part I: A Unified Formulation. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 28(1):26–37, 1998. Online available at <http://citeseer.ist.psu.edu/fonseca98multiobjective.html> [accessed 2007-07-29]. See also [34].
 - [16] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, July 1994. ISBN 0-2626-5040-1.
 - [17] Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, August 1991. ISBN 0-2620-6141-4, 978-0-262-06141-4.
 - [18] Vira Chankong and Yacov Y. Haimes. *Multiobjective Decision Making Theory and Methodology*. North-Holland, Elsevier, Dover Publications, New York, January 1983. ISBN 978-0444007100, 978-0486462899.
 - [19] Ralph E. Steuer. *Multiple Criteria Optimization: Theory, Computation and Application*. Krieger Pub Co, reprint edition, August 1989. ISBN 978-0894643934.
 - [20] Yacov Y. Haimes and Ralph E. Steuer, editors. *Proceedings of the 14th International Conference on Multiple Criteria Decision Making: Research and Practice in Multi Criteria Decision Making (MCDM'1998)*, volume 487 of *Lecture Notes in Economics and Mathematical Systems*, June 12–18, 1998, University of Virginia, Charlottesville, Virginia, USA. Springer. ISBN 978-3540672661. See <http://www.virginia.edu/~risk/mcdm98.html> [accessed 2007-09-10]. Published June 15, 2000.
 - [21] E. A. Galperin. Pareto analysis vis-à-vis balance space approach in multi-objective global optimization. *Journal of Optimization Theory and Applications*, 93(3):533–545, June 1997. ISSN 0022-3239 (Print) 1573-2878 (Online). doi:10.1023/A:1022639028824. Online available at <http://www.springerlink.com/content/m71638106736h581/fulltext.pdf> and <http://dx.doi.org/10.1023/A:1022639028824> [accessed 2007-09-10].
 - [22] Aharon Ben-Tal. Characterization of Pareto and lexicographic optimal solutions. In *Proceedings of the Third Conference on Multiple Criteria Decision Making: Theory and Application*, pages 1–11, 1979. In proceedings [368].
 - [23] T. Satoh, Tadashi Ishihara, and H. Inooka. Systematic design via the method of inequalities. *Control Systems Magazine, IEEE*, 16(5):57–65, October 1996. ISSN 0272-1708. doi:10.1109/37.537209.
 - [24] J. F. Whidborne, D.-W. Gu, and I. Postlethwaite. Algorithms for the Method of Inequalities – A Comparative Study. In *Proceedings of the 1995 American Control Conference*, pages 3393–3397, 1995, Seattle, Washington.
 - [25] P. Zhang and A.H. Coonick. Coordinated Synthesis of PSS Parameters in Multi-Machine Power Systems Using the Method of Inequalities Applied to Genetic Algorithms. *IEEE Transactions on Power Systems*, 15(2):811–816, May 2000.
 - [26] P. B. Wilson and M. D. Macleod. Low implementation cost IIR digital filter design using genetic algorithms. In *Proceedings of the IEE/IEEE Workshop on Natural Algorithms in Signal Processing*, pages 1–8, 1993, Chelmsford, UK.
 - [27] Abraham Charnes and William Wager Cooper. *Management Models and Industrial Applications of Linear Programming*. John Wiley & Sons Inc, New York, December 1961. ISBN 978-0471148500.

- [28] Kay Chen Tan, Eik Fun Khor, Tong Heng Lee, and Ramasubramanian Sathikannan from the National University of Singapore, 4 Engineering Drive 3, Singapore 117576, Republic of Singapore. An Evolutionary Algorithm with Advanced Goal and Priority Specification for Multi-objective Optimization. *Journal of Artificial Intelligence Research*, 18:183–215, February 2003. Online available at <http://www.jair.org/media/842/live-842-1969-jair.pdf> and <http://citeseer.ist.psu.edu/tan03evolutionary.html> [accessed 2007-08-25].
- [29] Kalyanmoy Deb from the Indian Institute of Technology Kanpur, Kanpur, India. Nonlinear goal programming using multi-objective genetic algorithms. *Journal of the Operational Research Society*, 52(3):291–302, March 2001. Online available at <http://www.lania.mx/~ccoello/EM00/deb01d.ps.gz> [accessed 2007-08-25].
- [30] Kalyanmoy Deb. Solving Goal Programming Problems Using Multi-Objective Genetic Algorithms. In *Proceedings of Congress on Evolutionary Computation*, volume 1, pages 77–84, 1999. In proceedings [502].
- [31] Ruhul A. Sarker, Hussein Abbas, and Samin Karim from the School of Computer Science, University of New South Wales, ADFA Campus, Northcott Drive, Canberra ACT, 2600, Australia. An Evolutionary Algorithm for Constrained Multiobjective Optimization Problems. In *Proceedings of the 5th Australia-Japan Joint Workshop on Intelligent & Evolutionary Systems (AJWIS'2001)*, pages 113–122, November 2001, The University of Otago, Dunedin, New Zealand. Online available at <http://www.lania.mx/~ccoello/EM00/sarker01a.pdf.gz> [accessed 2007-08-25].
- [32] Hartmut Pohlheim. GEATbx Introduction – Evolutionary Algorithms: Overview, Methods and Operators. Technical Report, <http://www.GEATbx.com> [accessed 2007-07-03], <http://www.GEATbx.com> [accessed 2007-07-03], November 2005. Documentation for GEATbx version 3.7 (Genetic and Evolutionary Algorithm Toolbox for use with Matlab). Online available at http://www.geatbx.com/download/GEATbx_Intro_Algorithmen_v38.pdf [accessed 2007-07-03].
- [33] Carlos M. Fonseca and Peter J. Fleming. An Overview of Evolutionary Algorithms in Multiobjective Optimization. *Evolutionary Computation*, 3(1):1–16, 1995. Online available at <http://citeseer.ist.psu.edu/108172.html> [accessed 2007-07-29].
- [34] Carlos M. Fonseca and Peter J. Fleming. Multiobjective optimization and multiple constraint handling with evolutionary algorithms – Part II: Application example. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 28(1):38–47, 1998. Online available at <http://citeseer.ist.psu.edu/27937.html> [accessed 2007-09-19]. See also [15].
- [35] Carlos M. Fonseca. Decision Making in Evolutionary Optimization (Abstract of Invited Talk). In *4th International Conference on Evolutionary Multi-Criterion Optimization*, 2007. In proceedings [518].
- [36] Carlos M. Fonseca and Peter J. Fleming. Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms. In *Practical Approaches to Multi-Objective Optimization*, 2004. In proceedings [509]. Online available at <http://drops.dagstuhl.de/opus/volltexte/2005/237/> [accessed 2007-09-19].
- [37] Carlos M. Fonseca. Preference Articulation in Evolutionary Multiobjective Optimisation – Plenary Talk. In *Proceedings of the Seventh International Conference on Hybrid Intelligent Systems, HIS 2007*, 2007. In proceedings [303].
- [38] Eik Fun Khor, Kay Chen Tan, Tong Heng Lee, and C. K. Goh. A Study on Distribution Preservation Mechanism in Evolutionary Multi-Objective Optimization. *Artificial Intelligence Review*, 23(1):31–33, March 2005. ISSN 0269-2821 (Print) 1573-7462 (Online). doi:10.1007/s10462-004-2902-3. Online available at <http://www.springerlink.com/content/h3w111g418gn1045/fulltext.pdf> [accessed 2007-08-25].
- [39] Hans Winkler. *Verbreitung und Ursache der Parthenogenesis im Pflanzen- und Tierreiche*. Verlag Gustav Fischer, Jena, 1920.

- [40] Joshua Lederberg and Alexa T. McCray. 'Ome Sweet 'Omics - A Genealogical Treasury of Words. *The Scientist*, 15(7):8, April 2001. Online available at <http://lhncbc.nlm.nih.gov/lhc/docs/published/2001/pub2001047.pdf> [accessed 2007-08-06].
- [41] Liviu Badea and Monica Stanciu. Refinement Operators Can Be (Weakly) Perfect. In *Inductive Logic Programming, Proceedings of 9th International Workshop, ILP-99 Bled, Slovenia*, volume 1634/1999 of *Lecture Notes in Computer Science (LNCS)*, pages 21–32. Springer Berlin/Heidelberg, June 24–27, 1999. ISBN 978-3-540-66109-2. doi:10.1007/3-540-48751-4_4. Online available at <http://www.ai.ici.ro/papers/ilp99.ps.gz> [accessed 2008-02-22].
- [42] Hendrik Skubch. Hierarchical Strategy Learning for FLUX Agents. Master's thesis, Technische Universität Dresden, Dresden, Germany, February 18, 2007. Supervisor: Prof. Michael Thielscher. Online available at <http://www.phenomene.de/hs/hslfa-dipl-hs.pdf> [accessed 2007-11-27]. See also [43].
- [43] Hendrik Skubch. *Hierarchical Strategy Learning for FLUX Agents*. An Applied Technique. VDM Verlag Dr. Müller, December 12, 2007. ISBN 978-3-8364-5271-7. See also [42].
- [44] Norman R. Paterson. *Genetic programming with context-sensitive grammars*. PhD thesis, Saint Andrew's University, September 2002. Online available at <ftp://ftp.dcs.st-and.ac.uk/pub/norman/GPwCSG.ps.gz> [accessed 2007-09-09].
- [45] Nicholas J. Radcliffe and Patrick D. Surry from the Department of Mathematics, University of Edinburgh, The Kings Buildings, EH9 3JZ, UK and Quadstone Ltd, 16 Chester Street, Edinburgh, EH3 7RA, UK. Fundamental limitations on search algorithms: Evolutionary computing in perspective. In Jan van Leeuwen, editor, *Computer Science Today – Recent Trends and Developments*, volume 1000 of *Lecture Notes in Computer Science (LNCS)*, pages 275–291. Springer-Verlag, 1995. ISBN 978-3-540-60105-0, 3-540-60105-8. doi:10.1007/BFb0015249. Online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.49.2376> [accessed 2008-08-12].
- [46] Elliott Sober. The Two Faces of Fitness. In Rama Shankar Singh, Costas B. Krimbas, Diane B. Paul, and John Beatty, editors, *Thinking about Evolution: Historical, Philosophical, and Political Perspectives*, chapter 15, pages 309–321. Cambridge University Press, 2001. ISBN 0521620708, 9780521620703. Online available at <http://philosophy.wisc.edu/sober/tff.pdf> [accessed 2008-08-11], Book partly online available at <http://books.google.de/books?id=HmVbYJ93d-AC> [accessed 2008-08-11].
- [47] Diane Paul. FITNESS: Historical Perspectives. In Evelyn Fox Keller and Elisabeth A. Lloyd, editors, *Keywords in Evolutionary Biology*, pages 112–114. Harvard University Press, Cambridge, MA, USA, November 1992. ISBN 978-0-674-50312-0, 0-674-50312-0.
- [48] Sewall Wright. The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution. In *Proceedings of the Sixth Annual Congress of Genetics*, volume 1, pages 356–366, 1932. Online available at <http://www.blackwellpublishing.com/ridley/classictexts/wright.pdf> [accessed 2007-08-11].
- [49] Stuart Alan Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, May 1993. ISBN 0195079515.
- [50] Sergey Gavrillets. *Fitness Landscapes and the Origin of Species*. Number MPB-41 in Monographs in Population Biology. Princeton University Press, July 2004. ISBN 069111983X.
- [51] Richard Dawkins. *Climbing Mount Improbable*. Penguin Books (UK) and W.W. Norton & Company (USA), London, 1 edition, 1996. ISBN 0670850187 and 0140179186, ASIN: 0393039307 and 0393316823.
- [52] Peter F. Stadler from the Institute for Theoretical Chemistry and Structural Biology, Währingerstraße 17, A-1090 Vienna, Austria. Fitness landscapes. In Michael Lässig and Angelo Valleriani, editors, *Biological Evolution and Statistical Physics*, Lecture Notes in Physics, pages 187–207. Springer-Verlag, Berlin, 2002. ISBN 3540431888,

- 978-3540431886. doi:10.1007/3-540-45692-9. Based on his talk *The Structure of Fitness Landscapes* given at the International Workshop on Biological Evolution and Statistical Physics, May 10–14, 2000, located at the Max-Planck-Institut für Physik komplexer Systeme, Nöthnitzer Str. 38, D-01187 Dresden, Germany. Online available at <http://www.bioinf.uni-leipzig.de/~studla/Publications/PREPRINTS/01-pfs-004.pdf> and <http://citeseer.ist.psu.edu/stadler02fitness.html> [accessed 2008-03-24].
- [53] K. Binder and A. P. Young. Spin glasses: Experimental facts, theoretical concepts, and open questions. *Reviews of Modern Physics*, 58(4):801–976, October 1986. doi:10.1103/RevModPhys.58.801. Online available at <http://link.aps.org/abstract/RMP/v58/p801> [accessed 2008-07-02].
 - [54] M. Mezard, Giorgio Parisi, and M. A. Virasoro. *Spin Glass Theory and Beyond – An Introduction to the Replica Method and Its Applications*, volume 9 of *World Scientific Lecture Notes in Physics*. World Scientific Publishing Company, November 1987. ISBN 978-9971-50-115-0, 9971-50-115-5, 978-9971-50-116-7, 9971-50-116-3.
 - [55] Melanie Mitchell. *An Introduction to Genetic Algorithms*. Complex Adaptive Systems. The MIT Press, reprint edition, February 1998. ISBN 0-2626-3185-7.
 - [56] Agoston E. Eiben and James E. Smith. *Introduction to Evolutionary Computing*. Natural Computing Series. Springer, first edition, November 2003. ISBN 978-3540401841. Partly online available at <http://books.google.de/books?id=7IOE5VIpFpwC> [accessed 2008-05-18].
 - [57] William B. Langdon and Riccardo Poli. *Foundations of Genetic Programming*. Springer, Berlin, first edition, January 2002. ISBN 3540424512.
 - [58] David H. Wolpert and William G. Macready. No Free Lunch Theorems for Search. Technical Report SFI-TR-95-02-010, The Santa Fe Institute, 1399 Hyde Park Rd., Santa Fe, NM, 87501, USA, February 6, 1995. Online available at <http://citeseer.ist.psu.edu/wolpert95no.html> and <http://www.santafe.edu/research/publications/workingpapers/95-02-010.pdf> [accessed 2008-03-28].
 - [59] David H. Wolpert and William G. Macready. No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, April 1997. doi:10.1109/4235.585893. Online available at <http://citeseer.ist.psu.edu/wolpert96no.html> [accessed 2008-03-28].
 - [60] Felix Streichert. Evolutionäre Algorithmen: Implementation und Anwendungen im Asset-Management-Bereich (Evolutionary Algorithms and their Application to Asset Management). Master’s thesis, Institut A für Mechanik, Universität Stuttgart, August 2001. Supervisors: Arnold Kirstner and Werner Koch. Online available at <http://www-ra.informatik.uni-tuebingen.de/mitarb/streiche/> [accessed 2007-08-17].
 - [61] Parag C. Pendharkar and Gary J. Koehler. A general steady state distribution based stopping criteria for finite length genetic algorithms. *European Journal of Operational Research*, 176(3):1436–1451, 2007. Online available at <http://linkinghub.elsevier.com/retrieve/pii/S0377221705008581> [accessed 2007-07-28].
 - [62] Karin Zielinski and Rainer Laur. Stopping Criteria for Constrained Optimization with Particle Swarms. In *Proceedings of the Second International Conference on Bioinspired Optimization Methods and their Application, BIOMA 2006*, pages 45–54, 2006. In proceedings [491] and (for an extended version) [63]. Online available at http://www.item.uni-bremen.de/staff/zilli/zielinski06stopping_PS0.pdf [accessed 2007-09-13].
 - [63] Karin Zielinski and Rainer Laur. Stopping Criteria for a Constrained Single-Objective Particle Swarm Optimization Algorithm. *Informatika*, 31(1):51–59, 2007. ISSN Print: 0350-5596, Web: 1854-387. Extended version of [62]. Online available at <http://www.item.uni-bremen.de/staff/zilli/zielinski07informatika.pdf> and <http://www.informatika.si/vols/vol31.html> [accessed 2007-09-13].

- [64] Method Man. Sub Crazy. In *Tical*, RZA, Method Man, and 4th Disciple, producers. Def Jam, November 15, 1994.
- [65] John R. Koza. *Genetic Programming, On the Programming of Computers by Means of Natural Selection*. A Bradford Book, The MIT Press, Cambridge, Massachusetts, 1992 first edition, 1993 second edition, 1992. ISBN 0262111705.
- [66] Anne E. Magurran. *Ecological Diversity and Its Measurement*. Princeton University Press, 41 William Street, Princeton, New Jersey 08540, USA, November 1988. ISBN 978-0-691-08491-6, 0-691-08485-8, 0691084912.
- [67] J. Shekel. Test Functions for Multimodal Search Techniques. In *Proceedings of the Fifth Annual Princeton Conference on Information Science and Systems*, 1971, Princeton, NJ, USA. See also http://en.wikipedia.org/wiki/Shekel_function [accessed 2007-11-06] and http://www-optima.amp.i.kyoto-u.ac.jp/member/student/hedar/Hedar_files/TestGO_files/Page2354.htm [accessed 2007-11-06].
- [68] Anatas Žilinskas from the Institute of Mathematics and Cybernetics, Academy of Science of the Lithuanian SSR, Vilnius, Lithuania. Algorithm AS 133: Optimization of One-Dimensional Multimodal Functions. *Applied Statistics*, 27(3):367–375, 1978. ISSN 00359254. doi:10.2307/2347182.
- [69] Rasmus K. Ursem. *Models for Evolutionary Algorithms and Their Applications in System Identification and Control Optimization*. PhD thesis, Department of Computer Science, University of Aarhus, Denmark, April 2003. Online available at http://www.daimi.au.dk/~ursem/publications/RKU_thesis_2003.pdf and <http://citeseer.ist.psu.edu/572321.html> [accessed 2007-08-14].
- [70] J. David Schaffer, Larry J. Eshelman, and Daniel Offutt. Spurious Correlations and Premature Convergence in Genetic Algorithms. In *Proceedings of the First Workshop on Foundations of Genetic Algorithms (FOGA)*, pages 102–112, June 1990. In proceedings [188].
- [71] Günter Rudolph. *Convergence Properties of Evolutionary Algorithms*, volume 35 of *Forschungsergebnisse zur Informatik*. Dr. Kovac, 1997, Hamburg, Germany. ISBN 978-3-86064-554-3, 3-86064-554-4.
- [72] Günter Rudolph. Self-adaptation and global convergence: A counter-example. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC99*, volume 1, pages 646–651, 1999. In proceedings [502]. Online available at <http://ls11-www.cs.uni-dortmund.de/people/rudolph/publications/papers/CEC99.pdf> and <http://citeseer.comp.nus.edu.sg/272402.html> [accessed 2008-02-24].
- [73] Günter Rudolph. Self-adaptive mutations may lead to premature convergence. *IEEE Transactions on Evolutionary Computation*, 5(4):410–414, 2001. See also [74].
- [74] Günter Rudolph. Self-adaptive mutations may lead to premature convergence. Technical Report CI-73/99, Fachbereich Informatik, Universität Dortmund, 44221 Dortmund, Germany, April 30, 2001. Online available at <http://citeseer.ist.psu.edu/444363.html> [accessed 2007-07-28]. See also [74].
- [75] William Michael Rudnick. *Genetic algorithms and fitness variance with an application to the automated design of artificial neural networks*. PhD thesis, Oregon Graduate Institute of Science & Technology, Beaverton, OR, USA, 1992. UMI Order No. GAX92-22642.
- [76] Dirk Thierens, David E. Goldberg, and Ângela Guimarães Pereira. Domino convergence, drift, and the temporal-salience structure of problems. In *Proceedings of the 1998 IEEE International Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence*, pages 535–540, 1998. doi:10.1109/ICEC.1998.700085. In proceedings [503]. Online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.46.6116> [accessed 2008-08-12].

- [77] Larry J. Eshelman and J. David Schaffer. Preventing Premature Convergence in Genetic Algorithms by Preventing Incest. In *ICGA, Proceedings of the 4th International Conference on Genetic Algorithms*, pages 115–122, 1991. In proceedings [749].
- [78] Shana Shiang-Fong Smith. Using multiple genetic operators to reduce premature convergence in genetic assembly planning. *Computers in Industry*, 54(1):35–49, 2004. ISSN 0166-3615.
- [79] Larry J. Eshelman, Richard A. Caruana, and J. David Schaffer. Biases in the crossover landscape. In *Proceedings of the third international conference on Genetic algorithms*, pages 10–19, 1989. In proceedings [652].
- [80] John Henry Holland. Genetic algorithms. *Scientific American*, 267(1):44–50, July 1992. Online available at <http://www.econ.iastate.edu/tesfatsi/holland.GAIntro.htm> and http://www.cc.gatech.edu/~turk/bio_sim/articles/genetic_algorithm.pdf [accessed 2007-08-29].
- [81] Agoston E. Eiben and C. A. Schippers. On Evolutionary Exploration and Exploitation. *Fundamenta Informaticae*, 35(1-4):35–50, 1998. Online available at <http://www.cs.vu.nl/~gusz/papers/FunInf98-Eiben-Schippers.ps> and <http://citeseer.ist.psu.edu/eiben98evolutionary.html> [accessed 2007-08-22].
- [82] Nitin Muttill and Shie-Yui Liong. Superior Exploration–Exploitation Balance in Shuffled Complex Evolution. *Journal of Hydraulic Engineering*, 130(12):1202–1205, December 2004.
- [83] Heni Ben Amor and Achim Rettinger. Intelligent exploration for genetic algorithms: Using self-organizing maps in evolutionary computation. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1531–1538, 2005. doi:10.1145/1068009.1068250. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068250> [accessed 2007-08-29].
- [84] Kalyanmoy Deb. Genetic Algorithms for Optimization. Technical Report, Kanpur Genetic Algorithms Laboratory (KanGAL), Kanpur, PIN 208 016, India, KanGAL, Kanpur, PIN 208 016, India, 2001. KanGAL Report Number 2001002, Online available at <http://citeseer.ist.psu.edu/604908.html> and <http://www.iitk.ac.in/kangal/papers/isna.ps.gz> [accessed 2007-07-28].
- [85] Christopher K. Oei, David E. Goldberg, and Shau-Jin Chang from the University of Illinois at Urbana-Champaign, Urbana, IL 61801. Tournament selection, niching, and the preservation of diversity. IlliGAL Report 91011, Illinois Genetic Algorithms Laboratory (IlliGAL), Department of Computer Science, Department of General Engineering, University of Illinois at Urbana-Champaign, December 1991. Online available at <http://www.illigal.uiuc.edu/pub/papers/IlliGALs/91011.ps.Z> [accessed 2008-03-15].
- [86] Simon Ronald. Preventing Diversity Loss in a Routing Genetic Algorithm with Hash Tagging. *Complexity International*, 2, 1995. ISSN 1320-0683. Online available at http://www.complexity.org.au/ci/vol02/sr_hash/ [accessed 2007-07-28].
- [87] Simon Ronald. *Genetic Algorithms and Permutation-Encoded Problems. Diversity Preservation and a Study of Multimodality*. PhD thesis, University Of South Australia. Department of Computer and Information Science, 1996.
- [88] Peter A. N. Bosman and Dirk Thierens from the Institute of Information and Computing Sciences, Utrecht University, P.O. Box 80.089, 3508, TB Utrecht, The Netherlands. Multi-objective optimization with diversity preserving mixture-based iterated density estimation evolutionary algorithms. *International Journal Approximate Reasoning*, 31(3):259–289, 2002. doi:10.1016/S0888-613X(02)00090-7. Received 1 January 2002; accepted 1 July 2002. Online available at [http://dx.doi.org/10.1016/S0888-613X\(02\)00090-7](http://dx.doi.org/10.1016/S0888-613X(02)00090-7) [accessed 2008-03-15].
- [89] Peter A. N. Bosman and Dirk Thierens. A Thorough Documentation of Obtained Results on Real-Valued Continuous And Combinatorial Multi-Objective Optimization Problems Using Diversity Preserving Mixture-Based Iterated Density Estima-

- tion Evolutionary Algorithms. Technical Report UU-CS-2002-052, Institute of Information and Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands, December 2002. Online available at <http://www.cs.uu.nl/research/techreps/repo/CS-2002/2002-052.pdf> [accessed 2007-08-24].
- [90] Marco Laumanns, Lothar Thiele, Kalyanmoy Deb, and Eckart Zitzler. On the Convergence and Diversity-Preservation Properties of Multi-Objective Evolutionary Algorithms. Technical Report 108, Computer Engineering and Networks Laboratory (TIK), Department of Electrical Engineering, Swiss Federal Institute of Technology (ETH) Zurich and Kanpur Genetic Algorithms Laboratory (KanGAL), Department of Mechanical Engineering, Indian Institute of Technology Kanpur, Gloriastrasse 35, CH-8092 Zurich, Switzerland, 2001. Online available at <http://e-collection.ethbib.ethz.ch/browse/alph/zitzlereckart.html> and <http://citeseer.ist.psu.edu/laumanns01convergence.html> [accessed 2007-08-14].
- [91] Sanaz Mostaghim. *Multi-objective Evolutionary Algorithms: Data structures, Convergence and, Diversity*. PhD thesis, Fakultät für Elektrotechnik, Informatik und Mathematik, Universität Paderborn, Deutschland (Germany), Shaker Verlag, 2004. Online available at <http://deposit.ddb.de/cgi-bin/dokserv?idn=974405604> and <http://ubdata.uni-paderborn.de/ediss/14/2004/mostaghi/disserta.pdf> [accessed 2007-08-17].
- [92] Christian Spieth, Felix Streichert, Nora Speer, and Andreas Zell from the Centre for Bioinformatics Tübingen (ZBIT), University of Tübingen, D-72076 Tübingen, Germany. Utilizing an Island Model for EA to Preserve Solution Diversity for Inferring Gene Regulatory Networks. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2004)*, volume 1, pages 146–151, 2004. In proceedings [497]. Online available at <http://www-ra.informatik.uni-tuebingen.de/software/JCell/publications.html> [accessed 2007-08-24].
- [93] Gulshan Singh and Kalyanmoy Deb. Comparison of multi-modal optimization algorithms based on evolutionary algorithms. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1305–1312, 2006. In proceedings [553]. Online available at <http://doi.acm.org/10.1145/1143997.1144200> [accessed 2007-09-10].
- [94] Steven Matt Gustafson. *An Analysis of Diversity in Genetic Programming*. PhD thesis, University of Nottingham, School of Computer Science & IT, Nottingham, U.K, February 2004. Online available at http://www.gustafsonresearch.com/thesis_html/index.html and <http://citeseer.ist.psu.edu/gustafson04analysis.html> [accessed 2008-07-22].
- [95] Steven Matt Gustafson, Anikó Ekárt, Edmund K. Burke, and Graham Kendall. Problem Difficulty and Code Growth in Genetic Programming. *Genetic Programming and Evolvable Machines*, 5(3):271–290, September 2004. ISSN 1389-2576 (Print) 1573-7632 (Online). doi:10.1023/B:GENP.0000030194.98244.e3. Online available at <http://www.gustafsonresearch.com/research/publications/gustafson-gpem2004.pdf> and <http://citeseer.ist.psu.edu/730858.html> [accessed 2008-07-21]. Submitted March 4, 2003; Revised August 7, 2003.
- [96] Edmund K. Burke, Steven Matt Gustafson, and Graham Kendall. Survey and analysis of diversity measures in genetic programming. In *Genetic and Evolutionary Computation Conference*, pages 716–723, 2002. In proceedings [563]. Online available at <http://www.gustafsonresearch.com/research/publications/gecco-diversity-2002.pdf> and <http://citeseer.ist.psu.edu/505881.html> [accessed 2008-07-22].
- [97] Edmund K. Burke, Steven Matt Gustafson, Graham Kendall, and Natalio Krasnogor. Is Increasing Diversity in Genetic Programming Beneficial? An Analysis of the Effects on Fitness. In *Congress on Evolutionary Computation*, pages 1398–1405, 2003. In proceedings [498]. Online available at <http://www.gustafsonresearch.com>

- com/research/publications/cec-2003.pdf and <http://citeseer.ist.psu.edu/700607.html> [accessed 2008-07-22].
- [98] Edmund K. Burke, Steven Matt Gustafson, and Graham Kendall. Diversity in genetic programming: An analysis of measures and correlation with fitness. *IEEE Transactions on Evolutionary Computation*, 8(1):47–62, February 2004. ISSN 1089-778X. doi:10.1109/TEVC.2003.819263. Online available at <http://www.gustafsonresearch.com/research/publications/gustafson-ieee2004.pdf> and <http://citeseer.ist.psu.edu/burke04diversity.html> [accessed 2008-07-22].
 - [99] Markus F. Brameier and Wolfgang Banzhaf. Explicit Control of Diversity and Effective Variation Distance in Linear Genetic Programming. In *EuroGP'02: Proceedings of the 5th European Conference on Genetic Programming*, pages 37–49, 2002. In proceedings [907]. Online available at <http://citeseer.ist.psu.edu/552561.html> [accessed 2007-09-09].
 - [100] Edwin D. de Jong, Richard A. Watson, and Jordan B. Pollack. Reducing Bloat and Promoting Diversity using Multi-Objective Methods. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 11–18, 2001. In proceedings [568]. Online available at <http://citeseer.ist.psu.edu/440305.html> and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/jong_2001_gecco.html [accessed 2008-07-20].
 - [101] Daniel N. Wilke, Schalk Kok, and Albert A. Groenwold. Comparison of linear and classical velocity update rules in particle swarm optimization: notes on diversity. *International Journal for Numerical Methods in Engineering*, 70(8):962–984, 2007. Online available at <http://doi.wiley.com/10.1002/nme.1867> [accessed 2007-08-20].
 - [102] Cage Kennylz. Radiohead. In *Radiohead / Agent Orange*, Necro, producer. Fondle 'Em Records, 1997. Also appears on the *For Your Box* mixtape.
 - [103] Stuart Alan Kauffman. Adaptation on rugged fitness landscapes. In Daniel L. Stein, editor, *Lectures in the Sciences of Complexity: The Proceedings of the 1988 Complex Systems Summer School*, volume Lecture I of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 527–618. Addison Wesley Publishing Company, Redwood City, June-July 1988, Santa Fe, New Mexico, USA. ISBN 978-0201510157, 0201510154. Published in September 1989.
 - [104] Aviv Bergman and Marcus W. Feldman from the Department of Biological Sciences, Stanford University, Stanford, California 94305, USA. Recombination dynamics and the fitness landscape. *Physica D: Nonlinear Phenomena*, 56(1):57–67, April 1992. ISSN 0167-2789. doi:10.1016/0167-2789(92)90050-W. Online available at [http://dx.doi.org/10.1016/0167-2789\(92\)90050-W](http://dx.doi.org/10.1016/0167-2789(92)90050-W) [accessed 2008-07-20].
 - [105] Ingo Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog Verlag, Stuttgart, 1973. ISBN 978-3772803741. his dissertation from 1970.
 - [106] Ingo Rechenberg. *Evolutionsstrategie '94*, volume 1 of *Werkstatt Bionik und Evolutionstechnik*. Frommann Holzboog, Stuttgart, September 1994. ISBN 9783772816420.
 - [107] Justinian P. Rosca and Dana H. Ballard from the Computer Science Department, University of Rochester, Rochester NY 14627, USA. Causality in Genetic Programming. In *Proceedings of the 6th International Conference on Genetic Algorithms (ICGA95)*, pages 256–263, 1995. In proceedings [747]. Online available at <http://citeseer.ist.psu.edu/rosca95causality.html> [accessed 2007-08-12].
 - [108] Stefan Droste and Dirk Wiesmann. On Representation and Genetic Operators in Evolutionary Algorithms. Technical Report CI-41/98, Fachbereich Informatik, Universität Dortmund, 44221 Dortmund, June 1998. Online available at <http://citeseer.ist.psu.edu/323494.html> and <http://en.scientificcommons.org/336032> [accessed 2007-08-12].

- [109] Reinhard Lohmann. Structure evolution and neural systems. In *Dynamic, Genetic, and Chaotic Programming: The Sixth-Generation*, pages 395–411. Wiley-Interscience, 1992. In collection [1912].
- [110] Reinhard Lohmann. Structure evolution and incomplete induction. *Biological Cybernetics*, 69(4):319–326, August 1993. ISSN 0340-1200 (Print) 1432-0770 (Online). doi:10.1007/BF00203128. Online available at <http://www.springerlink.com/content/q2q668316m771073/fulltext.pdf> [accessed 2007-08-12].
- [111] Justinian P. Rosca. An Analysis of Hierarchical Genetic Programming. Technical Report TR566, The University of Rochester, Computer Science Department, Rochester, New York, 1527, USA, March 1995. Online available at <ftp://ftp.cs.rochester.edu/pub/u/rosca/gp/95.tr566.ps.gz> and <http://citeseer.ist.psu.edu/rosca95analysis.html> [accessed 2008-02-24].
- [112] Christian Igel. Causality of Hierarchical Variable Length Representations. In *Proceedings of the 1998 IEEE World Congress on Computational Intelligence*, pages 324–329, 1998. IEEE Press, Anchorage, Alaska, USA. Online available at <http://www.neuroinformatik.ruhr-uni-bochum.de/PEOPLE/igel/CoHVLr.ps.gz> and <http://citeseer.ist.psu.edu/61421.html> [accessed 2007-08-12].
- [113] Bernhard Sendhoff, Martin Kreutz, and Werner von Seelen. A Condition for the Genotype-Phenotype Mapping: Causality. In *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA97)*, 1997. In proceedings [746]. Online available at <http://citeseer.ist.psu.edu/sendhoff97condition.html> and <http://arxiv.org/abs/adap-org/9711001> [accessed 2007-08-12].
- [114] Peter Stagge and Christian Igel. Structure Optimization and Isomorphisms. In *Theoretical Aspects of Evolutionary Computing*, pages 409–422. Springer, 2000. In collection [615]. Online available at <http://citeseer.ist.psu.edu/520775.html> [accessed 2007-08-13].
- [115] Franz Rothlauf. *Representations for Genetic and Evolutionary Algorithms*. Physica-Verlag, second edition, August 2002 (1st ed.), 2006 (2nd ed.). ISBN 978-3790814965, 354025059-X, 978-354025059-3. Foreword by David E. Goldberg. Partly online available at <http://books.google.de/books?id=fQrSUwop4JkC> [accessed 2008-02-27].
- [116] Krasimir Kolarov from the Interval Research Corporation, 1801-C Page Mill Road, Palo Alto, CA 94304, USA. Landscape Ruggedness in Evolutionary Algorithms. In *Proceedings of The IEEE Conference on Evolutionary Computation*, pages 19–24, 1997. In proceedings [504]. Online available at <http://citeseer.ist.psu.edu/14650.html> and <http://de.scientificcommons.org/27191> [accessed 2007-08-19].
- [117] David C. Wedge and Douglas B. Kell from the Manchester Interdisciplinary Biocentre, School of Chemistry, The University of Manchester, 131 Princess Street, Manchester, M1 7DN, United Kingdom. Rapid Prediction of Optimum Population Size in Genetic Programming using a Novel Genotype-Fitness Correlation. In *Genetic and Evolutionary Computation Conference*, 2008. In proceedings [550].
- [118] Lee Altenberg from the Hawaii Institute of Geophysics and Planetology, University of Hawaii at Manoa, Honolulu, HI USA 96822. Fitness Distance Correlation Analysis: An Instructive Counterexample. In *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA97)*, pages 57–64, 1997. In proceedings [746]. Online available at <http://dynamics.org/Altenberg/PAPERS/FDCAIIC/> and <http://www.cs.uml.edu/~giam/91.510/Papers/Altenberg1997.pdf> [accessed 2008-07-20].
- [119] Edward D. Weinberger. Correlated and uncorrelated Fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63:325–336, 1990.
- [120] Bernard Manderick, Mark de Weger, and Piet Spiessens. The Genetic Algorithm and the Structure of the Fitness Landscape. In *Proceedings of the 4th International Conference on Genetic Algorithms*, pages 143–150, 1991. In proceedings [749].

- [121] Terry Jones and Stephanie Forrest. Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 184–192, 1995. In proceedings [747]. Online available at <http://citeseer.ist.psu.edu/jones95fitness.html> and <http://www.santafe.edu/research/publications/workingpapers/95-02-022.ps> [accessed 2008-07-20].
- [122] Terry Jones. *Evolutionary Algorithms, Fitness Landscapes and Search*. PhD thesis, The University of New Mexico. Albuquerque, New Mexico, USA, May 1995. Online available at <http://www.cs.unm.edu/~forrest/dissertations-and-proposals/terry.pdf> [accessed 2008-07-02].
- [123] Manuel Clergue, Philippe Collard, Marco Tomassini, and Leonardo Vanneschi. Fitness Distance Correlation And Problem Difficulty For Genetic Programming. In *GECCO'02: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 724–732, 2002. In proceedings [563]. Online available at <http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/gecco2002/gecco-2002-14.pdf> and <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2002/GP072.pdf> [accessed 2008-07-20]. See also [124].
- [124] Leonardo Vanneschi and Marco Tomassini from the Computer Science Institute, University of Lausanne, Switzerland. A Study on Fitness Distance Correlation and Problem Difficulty for Genetic Programming. In *Graduate Student Workshop, GECCO Conference*, pages 307–310, 2002. In proceedings [567]. Online available at http://personal.disco.unimib.it/Vanneschi/GECCO_2002_PHD_WORKSHOP.pdf and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/vanneschi_2002_gecco_workshop.html [accessed 2008-07-20]. See also [123].
- [125] Hans-Georg Beyer. Toward a Theory of Evolution Strategies: The (μ, λ) -Theory. *Evolutionary Computation*, 2(4):381–407, 1994. Online available at <http://ls11-www.cs.uni-dortmund.de/people/beyer/coll/Bey95a/Bey95a.ps> and <http://citeseer.ist.psu.edu/249681.html> [accessed 2007-08-27].
- [126] Lee Altenberg. The Schema Theorem and Price's Theorem. In *Foundations of Genetic Algorithms 3*, pages 23–49, 1994. In proceedings [737]. Online available at <http://dynamics.org/Altenberg/PAPERS/STPT/> and <http://citeseer.ist.psu.edu/altenberg95schema.html> [accessed 2008-07-20].
- [127] John J. Grefenstette from the Navy Center for Artificial Intelligence, Code 5514, Naval Research Laboratory, Washington, DC 20375-5337, USA. Predictive Models Using Fitness Distributions of Genetic Operators. In *Foundations of Genetic Algorithms 3*, pages 139–161, 1994. In proceedings [737]. Online available at http://reference.kfupm.edu.sa/content/p/r/predictive_models_using_fitness_distribu_672029.pdf and <http://citeseer.ist.psu.edu/grefenstette95predictive.html> [accessed 2008-07-20].
- [128] Nicholas J. Radcliffe and Patrick David Surry. Fitness variance of formae and performance prediction. In *Foundations of Genetic Algorithms 3*, pages 51–72, 1994. In proceedings [737]. Online available at <http://citeseer.comp.nus.edu.sg/radcliffe94fitness.html> [accessed 2007-07-28].
- [129] Colin R. Reeves and Christine Wright from the Statistics and Operational Research Division, School of Mathematical and Information Sciences, Coventry University, UK. An Experimental Design Perspective on Genetic Algorithms. In *Foundations of Genetic Algorithms 3*, pages 7–22, 1994. In proceedings [737]. Online available at <http://citeseer.ist.psu.edu/reeves95experimental.html> [accessed 2008-07-20].
- [130] Manfred Eigen and Peter Schuster. *The Hypercycle: A Principle of Natural Self Organization*. Springer-Verlag, June 1979. ISBN 0387092935, 978-0387092935.
- [131] M. Nowak and P. Schuster from the Institut für Theoretische Chemie, Universität Wien, Austria. Error thresholds of replication in finite populations mutation fre-

- quencies and the onset of Muller's ratchet. *Journal of Theoretical Biology*, 137(4): 375–395, April 20, 1989. ISSN 0022-5193.
- [132] Gabriela Ochoa, Inman Harvey, and Hilary Buxton from the Centre for the Study of Evolution, Centre for Computational Neuroscience and Robotics, COGS – the University of Sussex, Falmer, Brighton BN1 9QH, UK. Error Thresholds and Their Relation to Optimal Mutation Rates. In *European Conference on Artificial Life*, pages 54–63, 1999. In proceedings [1913]. Online available at <http://citeseer.ist.psu.edu/ochoa99error.html> and <ftp://ftp.cogs.susx.ac.uk/pub/users/inmanh/ecal99.ps.gz> [accessed 2008-07-20].
- [133] Leonardo Vanneschi, Manuel Clergue, Philippe Collard, Marco Tomassini, and Sébastien Vérel. Fitness Clouds and Problem Hardness in Genetic Programming. In *Genetic and Evolutionary Computation – GECCO-2004, Part II*, pages 690–701, 2004. doi:10.1007/b98645. In proceedings [559]. Online available at <http://www.i3s.unice.fr/~verel/publi/gecco04-FCandPbHardnessGP.pdf> and <http://hal.archives-ouvertes.fr/hal-00160055/en/> [accessed 2007-10-14], online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/vanneschi_fca_gecco2004.html and http://hal.inria.fr/docs/00/16/00/55/PDF/fc_gp.pdf [accessed 2008-07-20].
- [134] Leonardo Vanneschi, Marco Tomassini, Philippe Collard, and Sébastien Vérel. Negative Slope Coefficient. A Measure to Characterize Genetic Programming. In *Proceedings of the 9th European Conference on Genetic Programming*, pages 178–189, 2006. In proceedings [903]. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/eurogp06_VanneschiTomassiniCollardVerel.html [accessed 2008-07-20].
- [135] Yuval Davidor from the Department of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel. Epistasis Variance: Suitability of a Representation to Genetic Algorithms. *Complex Systems*, 4(4):369–383, 1990. ISSN 0891-2513.
- [136] Sébastien Verel, Philippe Collard, and Manuel Clergue. Measuring the evolvability landscape to study neutrality. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 613–614, 2006. In proceedings [553]. Online available at <http://portal.acm.org/citation.cfm?id=1144107> and <http://arxiv.org/abs/0709.4011> [accessed 2008-02-27].
- [137] The RZA. Pit Full of Snakes. In Gravediggaz, editor, *The Pick, the Sickie and the Shovel*, 4th Disciple and True Master, producers. Gee Street/V2/BMG Records, October 14, 1997.
- [138] Albert Donally Bethke. *Genetic algorithms as function optimizers*. PhD thesis, University of Michigan, University of Michigan, Ann Arbor, MI, USA, 1980. Order No. AAI8106101, Dissertation Abstracts International, 41(9), 3503B (University Microfilms No. 8106101). See also <http://hdl.handle.net/2027.42/3572> [accessed 2007-10-28].
- [139] David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, first edition, January 1989. ISBN 0201157675.
- [140] Gunar E. Liepins and Michael D. Vose. Deceptiveness and Genetic Algorithm Dynamics. In *Proceedings of the First Workshop on Foundations of Genetic Algorithms (FOGA)*, pages 36–50, 1991. In proceedings [188]. Online available at <http://www.osti.gov/bridge/servlets/purl/6445602-CfqU6M/> [accessed 2007-11-05].
- [141] Satoshi Watanabe. *Knowing and Guessing: A Quantitative Study of Inference and Information*. John Wiley & Sons, New York, USA, 1969. ISBN 0471921300, 9780471921301.
- [142] Tom M. Mitchell from the Computer Science Department, Rutgers University, New Brunswick, NJ 08903, U.S.A. Generalization as Search. In Bonnie Lynn Webber and Nils J. Nilsson, editors, *Readings in Artificial Intelligence*, pages 517–542. Tioga Pub. Co. Press / Morgan Kaufmann Publishers / Elsevier Science & Technology

- Books, 2nd edition, February 1981. ISBN 0935382038, 978-0935382037, 0934613036, 978-0934613033. See also [143].
- [143] Tom M. Mitchell from the Computer Science Department, Rutgers University, New Brunswick, NJ 08903, U.S.A. Generalization as Search. *Artificial Intelligence*, 18(2): 203–226, March 1982. ISSN 0004-3702. doi:10.1016/0004-3702(82)90040-6. See also [142].
 - [144] Dirk Thierens from the Department of Computer Science, Utrecht University, P.O. Box 80089, 3508 TB Utrecht, The Netherlands. On the scalability of simple genetic algorithms. Technical Report UU-CS-1999-48, Department of Information and Computing Sciences, Utrecht University, 1999. Persistent Identifier: URN:NBN:NL:UI:10-1874/18986. Online available at http://igitur-archive.library.uu.nl/math/2007-0118-201740/thierens_99_on_the_scalability.pdf and <http://www.cs.uu.nl/research/techreps/repo/CS-1999/1999-48.pdf> [accessed 2008-08-12], , online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.40.5295> [accessed 2008-08-12].
 - [145] Lionel Barnett from the Centre for Computational Neuroscience and Robotics, Centre for the Study of Evolution, Department of Cognitive and Computing Sciences, University of Sussex, Brighton BN1 9QH, UK. Ruggedness and neutrality – the NKp family of fitness landscapes. In *Artificial Life VI: Proceedings of the sixth international conference on Artificial life*, pages 18–27, 1998. In proceedings [1914]. Online available at ftp://ftp.informatics.sussex.ac.uk/pub/users/lionelb/publications/alife6_paper.pdf and <http://citeseer.ist.psu.edu/barnett98ruggedness.html> [accessed 2008-06-01], online available at <http://citeseer.ist.psu.edu/144047.html> [accessed 2008-02-27].
 - [146] Richard Dawkins. The Evolution of Evolvability. In *ALIFE – Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, pages 201–220, 1987. In proceedings [1558]. Published in January 1989.
 - [147] Andreas Wagner. *Robustness and Evolvability in Living Systems*. Princeton Studies in Complexity. Princeton University Press, August 2005. ISBN 0691122407.
 - [148] Andreas Wagner. Robustness, evolvability, and neutrality. *FEBS Lett*, 579(8):1772–1778, March 21, 2005. ISSN 0014-5793. Edited by Robert Russell and Giulio Superti-Furga. Online available at <http://www.citeulike.org/user/timflutre/article/297788> and <http://dx.doi.org/10.1016/j.febslet.2005.01.063> [accessed 2007-08-05].
 - [149] Marc Kirschner and John Gerhart. Evolvability. *Proceedings of the National Academy of Science of the USA (PNAS)*, 95(15):8420–8427, July 21, 1998. ISSN 0027-8424. Online available at <http://www.pnas.org/cgi/content/full/95/15/8420> [accessed 2007-08-05].
 - [150] Rupert J. Riedl. A systems-analytical approach to macroevolutionary phenomena. *Quarterly Review of Biology*, (52):351–370, 1977.
 - [151] Lee Altenberg. Genome Growth and the Evolution of the Genotype-Phenotype Map. In *Evolution and Biocomputation – Computational Models of Evolution*, pages 205–259. Springer-Verlag, 1995. In collection [605]. Online available at <http://citeseer.ist.psu.edu/254722.html> [accessed 2007-07-29]; corrections and revisions, October 1997.
 - [152] Günter P. Wagner and Lee Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, June 1996. ISSN 00143820. Online available at <http://citeseer.ist.psu.edu/437225.html> and <http://dynamics.org/Altenberg/PAPERS/CAEE/> [accessed 2007-08-05].
 - [153] John Tyler Bonner. *On Development: The Biology of Form*. Commonwealth Fund Publications. Harvard University Press, Cambridge, MA, USA, new ed edition, December 12, 1974. ISBN 0674634128, 978-0674634121.
 - [154] William Beaudoin, Sébastien Verel, Philippe Collard, and Cathy Esczut from the University of Nice-Sophia Antipolis, I3S Laboratory, Sophia Antipolis, France. Deceptiveness and neutrality the ND family of fitness landscapes. In *GECCO'06: Pro-*

- ceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 507–514, 2006. doi:10.1145/1143997.1144091. In proceedings [553]. Online available at <http://doi.acm.org/10.1145/1143997.1144091> [accessed 2007-08-25].
- [155] Christian V. Forst, Christian Reidys, and Jacqueline Weber from the Institut für Molekulare Biotechnologie, Beutenbergstraße 11, PF 100 813, D-07708 Jena, Germany. Evolutionary Dynamics and Optimization: Neutral Networks as Model-Landscapes for RNA Secondary-Structure Folding-Landscapes. In *European Conference on Artificial Life*, pages 128–147, 1995. In proceedings [1915]. Online available at <http://arxiv.org/abs/adap-org/9505002v1> and <http://citeseer.ist.psu.edu/8900.html> [accessed 2008-06-01].
- [156] Marijn A. Huynen, Peter F. Stadler, and Walter Fontana. Smoothness within ruggedness: The role of neutrality in adaptation. *Proceedings of the National Academy of Science, USA*, 93:397–401, January 1996. Evolution. Online available at <http://www.pnas.org/cgi/reprint/93/1/397.pdf> and <http://citeseer.ist.psu.edu/91710.html> [accessed 2008-02-27]. Communicated by Hans Frauenfelder, Los Alamos National Laboratory, Los Alamos, NM, September 20, 1995 (received for review June 29, 1995).
- [157] Marc Toussaint and Christian Igel. Neutrality: A Necessity for Self-Adaptation. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2002)*, pages 1354–1359, 2002. In proceedings [499]. Online available at <http://citeseer.ist.psu.edu/toussaint02neutrality.html> [accessed 2007-07-28].
- [158] Nils Eldredge and Stephen Jay Gould. Punctuated equilibria: an alternative to phyletic gradualism. In Thomas J. M. Schopf, editor, *Models in Paleobiology*, chapter 5, pages 82–115. Freeman, Cooper & Co / DoubleDay, San Francisco, USA, January 1972. ISBN 038503024X, 978-0385030243. Online available at <http://www.blackwellpublishing.com/ridley/classictexts/eldredge.pdf> [accessed 2008-07-01].
- [159] James P. Cohoon, S. U. Hegde, Worthy Neil Martin, and D. Richards. Punctuated equilibria: a parallel genetic algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 148–154, 1987. In proceedings [750].
- [160] Worthy Neil Martin, Jens Lienig, and James P. Cohoon. Island (migration) Models: Evolutionary Algorithms Based on Punctuated Equilibria. In *Handbook of Evolutionary Computation*, chapter 6.3. Oxford University Press, 1997. In collection [410]. Online available at http://www.cs.virginia.edu/papers/Island_Migration.pdf [accessed 2007-08-13].
- [161] Erik van Nimwegen, James P. Crutchfield, and Melanie Mitchell from the Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA. Statistical Dynamics of the Royal Road Genetic Algorithm. *Theoretical Computer Science*, 229(1–2):41–102, November 6, 1999. ISSN 0304-3975. doi:10.1016/S0304-3975(99)00119-X. Online available at [http://dx.doi.org/10.1016/S0304-3975\(99\)00119-X](http://dx.doi.org/10.1016/S0304-3975(99)00119-X) and <http://citeseer.ist.psu.edu/nimwegen98statistical.html> [accessed 2008-06-01].
- [162] Tom Smith, Phil Husbands, Paul Layzell, and Michael O’Shea. Fitness landscapes and evolvability. *Evolutionary Computation*, 10(1):1–34, Spring 2002. ISSN 1063-6560. doi:10.1162/106365602317301754. Online available at http://wotan.liu.edu/docis/dbl/evocom/2002_10_1_1_FLAЕ.html [accessed 2007-07-28].
- [163] Wolfgang Banzhaf, Guillaume Beslon, Steffen Christensen, James A. Foster, François Képès, Virginie Lefort, Julian Francis Miller, Miroslav Radman, and Jeremy J. Ramsden. From artificial evolution to computational evolution: a research agenda. *Nature Reviews Genetics*, 7(9):729–735, September 2006. ISSN 1471-0056, 1471-0064. doi:10.1038/nrg1921. Guidelines/Perspectives. Online available at <http://www.cs.mun.ca/~banzhaf/papers/nrg2006.pdf> [accessed 2008-07-20].
- [164] Ting Hu and Wolfgang Banzhaf from the Department of Computer Science, Memorial University of Newfoundland, St. John’s, NL, A1B 3X5, Canada. Measuring Rate of

- Evolution in Genetic Programming Using Amino Acid to Synonymous Substitution Ratio k_a/k_s . In *Genetic and Evolutionary Computation Conference*, pages 1337–1338, 2008. In proceedings [550]. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/Hu_2008_gecco.html and <http://www.cs.mun.ca/~tingh/papers/GECCO08.pdf> [accessed 2008-08-08].
- [165] Ziheng Yang and Joseph P. Bielawski. Statistical methods for detecting molecular adaptation. *Trends in Ecology & Evolution*, 15(12):496–503, December 1, 2000. ISSN 0169-5347. doi:10.1016/S0169-5347(00)01994-7. Online available at [http://dx.doi.org/10.1016/S0169-5347\(00\)01994-7](http://dx.doi.org/10.1016/S0169-5347(00)01994-7) and <http://citeseer.ist.psu.edu/yang00statistical.html> [accessed 2008-07-20].
- [166] Robert P. Worden from the Logica UK Ltd, Betjeman House, 104 Hills Road, Cambridge CB2 1LQ, U.K. A Speed Limit for Evolution. *Journal of Theoretical Biology*, 176(1):137–152, September 7, 1995. ISSN 0022-5193. doi:10.1006/jtbi.1995.0183. Online available at <http://dx.doi.org/10.1006/jtbi.1995.0183> [accessed 2008-08-10].
- [167] Martijn A. Huynen. Exploring Phenotype Space Through Neutral Evolution. *Journal of Molecular Evolution*, 43(3):165–169, September 1996. ISSN 0022-2844 (Print) 1432-1432 (Online). doi:10.1007/PL00006074. Online available at <http://citeseer.ist.psu.edu/32668.html> and <http://www.santafe.edu/research/publications/wpabstract/199510100> [accessed 2008-02-27].
- [168] Terry Stewart from the Carleton University, Institute of Interdisciplinary Studies (Cognitive Science), Ottawa, Canada. Extrema selection: accelerated evolution on neutral networks. In *Congress on Evolutionary Computation*, volume 1, 2001. doi:10.1109/CEC.2001.934366. In proceedings [500]. Online available at http://rob.ccmlab.ca:8080/~terry/papers/2001-Extrema_Selection.pdf and <http://citeseer.ist.psu.edu/654440.html> [accessed 2008-06-01].
- [169] Lionel Barnett. Tangled Webs: Evolutionary Dynamics on Fitness Landscapes with Neutrality. Master's thesis, School of Cognitive Science, University of East Sussex, Brighton, Brighton, 1997. Supervisor: Inman Harvey. Online available at <http://citeseer.ist.psu.edu/barnett97tangled.html> and <ftp://ftp.cogs.susx.ac.uk/pub/users/inmanh/lionelb/> [accessed 2007-08-13].
- [170] Erich Bornberg-Bauer and Hue Sun Chan. Modeling evolutionary landscapes: Mutational stability, topology, and superfunnels in sequence space. *Proceedings of the National Academy of Science of the United States of America (PNAS) – Biophysics*, 96(19):10689–10694, September 14, 1999. Edited by Peter G. Wolynes. Online available at <http://www.pnas.org/cgi/content/abstract/96/19/10689> [accessed 2008-07-02].
- [171] Erik van Nimwegen, James P. Crutchfield, and Martijn Huynen. Neutral evolution of mutational robustness. *Proceedings of the National Academy of Science of the United States of America (PNAS) – Evolution*, 96(17):9716–9720, August 17, 1999. Online available at <http://www.pnas.org/cgi/reprint/96/17/9716.pdf> [accessed 2008-07-02].
- [172] Erik van Nimwegen and James P. Crutchfield from the Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA. Optimizing Epochal Evolutionary Search: Population-Size Dependent Theory. *Machine Learning*, 45(1):77–114, October 2001. ISSN 0885-6125 (Print) 1573-0565 (Online). doi:10.1023/A:1012497308906. Editor: Leslie Pack Kaelbling. Online available at <http://www.springerlink.com/content/nv23534338656352/fulltext.pdf> [accessed 2008-07-02].
- [173] Claus O. Wilke from the Digital Life Lab, Mail-Code 136-93, California Institute of Technology, Pasadena, CA 91125, U.S.A. Adaptive Evolution on Neutral Networks. *Bulletin of Mathematical Biology*, 63(4):715–730, July 2001. ISSN 0092-8240 (Print) 1522-9602 (Online). Online available at <http://arxiv.org/abs/physics/0101021v1> [accessed 2008-07-02].
- [174] Rob Shipman, Marc Shackleton, Marc Ebner, and R. Watson. Neutral search spaces for artificial evolution: a lesson from life. In Mark Bedau, John S. McCaskill, Norman H. Packard, Steen Rasmussen, John McCaskill, and Norman Packard, editors,

- Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*, Bradford Books, Complex Adaptive Systems. The MIT Press, July 2000. ISBN 026252290X, 978-0262522908. Online available at <http://citeseer.ist.psu.edu/shipman00neutral.html> and <http://books.google.com/books?id=-xLGm7KFGy4C> [accessed 2007-07-29].
- [175] Mark Shackleton, Rob Shipman, and Marc Ebner. An investigation of redundant genotype-phenotype mappings and their role in evolutionary search. In *Proceedings of the 2000 Congress on Evolutionary Computation CEC00*, pages 493–500, 2000. In proceedings [501]. Online available at <http://citeseer.ist.psu.edu/409243.html> [accessed 2007-07-29].
- [176] Rob Shipman, Mark Shackleton, and I. Harvey. The use of neutral genotype-phenotype Mappings for improved evolutionary search. *BT Technology Journal*, 18(4):103–111, 2000. Online available at <http://citeseer.ist.psu.edu/shipman00use.html> [accessed 2007-07-29].
- [177] Rob Shipman from the Future Technologies Group, Complex Systems Laboratory. Genetic Redundancy: Desirable or problematic for evolutionary adaptation? In *Proceedings of the 4th International Conference on Artificial Neural Nets and Genetic Algorithms*, pages 1–11, 1999. In proceedings [579]. Online available at <http://citeseer.ist.psu.edu/shipman99genetic.html> [accessed 2007-07-29].
- [178] Vesselin K. Vassilev and Julian Francis Miller. The Advantages of Landscape Neutrality in Digital Circuit Evolution. In *ICES'00: Proceedings of the Third International Conference on Evolvable Systems*, volume 1801/2000, pages 252–263. Springer-Verlag, London, UK, April 17-19, 2000, Edinburgh, Scotland, UK. ISBN 3-540-67338-5. Online available at <http://citeseer.ist.psu.edu/vassilev00advantages.html> [accessed 2007-11-02].
- [179] Gwoing Tina Yu and Julian Francis Miller. Finding Needles in Haystacks Is Not Hard with Neutrality. In *EuroGP'02: Proceedings of the 5th European Conference on Genetic Programming*, pages 13–25, 2002. In proceedings [907]. Online available at <http://citeseer.ist.psu.edu/yu02finding.html> and http://www.cs.mun.ca/~tinayu/index_files/addr/public_html/EuroGP2002.pdf [accessed 2007-11-02].
- [180] Patrick C. Phillips from the Biology Department, University of Texas, Arlington, Texas 76019-0498, USA. The Language of Gene Interaction. *Genetics*, 149(3):1167–1171, July 7, 1998. ISSN 0016-6731. Online available at <http://www.genetics.org/cgi/content/full/149/3/1167> [accessed 2008-03-02].
- [181] William Bateson. *Mendel's Principles of Heredity*. Cambridge University Press, Cambridge, 1909. ISBN 9781428648197. 1930: fourth impression of the 1909 edition.
- [182] Sir Ronald Aylmer Fisher. The correlations between relatives on the supposition of Mendelian inheritance. *Philosophical Transactions of the Royal Society of Edinburgh*, 52:399–433, 1918. See also http://en.wikipedia.org/wiki/The_Correlation_Between_Relatives_on_the_Supposition_of_Mendelian_Inheritance [accessed 2008-03-02].
- [183] Jay L. Lush from the Iowa State College, Ames, Iowa. Progeny test and individual performance as indicators of an animal's breeding value. *Journal of Dairy Science*, 18(1):1–19, January 1935. Online available at <http://jds.fass.org/cgi/reprint/18/1/1> [accessed 2007-11-27].
- [184] George C. Williams. Pleiotropy, Natural Selection, and the Evolution of Senescence. *Evolution*, 11(4):398–411, December 1957. doi:10.2307/2406060. See also [1916].
- [185] Yuval Davidor. Epistasis Variance: A Viewpoint on GA-Hardness. In *Proceedings of the First Workshop on Foundations of Genetic Algorithms*, pages 23–35, 1990. In proceedings [188].
- [186] Lee Altenberg from the Hawai'i Institute of Geophysics and Planetology, University of Hawaii at Manoa, Honolulu, HI USA. NK Fitness Landscapes. In *Handbook of Evolutionary Computation*, chapter B2.7.2. Oxford University Press, November 27, 1996.

- In collection [410]. Online available at <http://citeseer.ist.psu.edu/704814.html> and <http://dynamics.org/Altenberg/FILES/LeeNKFL.pdf> [accessed 2007-11-27].
- [187] Bart Naudts and Alain Verschoren. Epistasis On Finite And Infinite Spaces. In *Proceedings of the 8th International Conference on Systems Research, Informatics and Cybernetics*, pages 19–23, 1996. Online available at <http://en.scientificcommons.org/155291> and <http://citeseer.ist.psu.edu/142750.html> [accessed 2007-08-13].
- [188] Bruce M. Spatz and Gregory J. E. Rawlins, editors. *Proceedings of the First Workshop on Foundations of Genetic Algorithms*, July 15–18, 1990, Indiana University, Bloomington Campus, Indiana, USA. Morgan Kaufmann Publishers, Inc., 2929 Campus Drive, Suite 260, San Mateo, CA 94403, USA. ISBN 1-55860-170-8. Published July 1, 1991. Partly online available at <http://books.google.de/books?id=Df12yLrlUZYC> [accessed 2008-05-29].
- [189] Bart Naudts, Dominique Suys, and Alain Verschoren. Generalized Royal Road Functions and their Epistasis. *Computers and Artificial Intelligence*, 19(4), 2000. Original: March 5, 1997. Online available at <http://citeseer.ist.psu.edu/111356.html> [accessed 2007-08-13].
- [190] Bart Naudts and Alain Verschoren. Epistasis and Deceptivity. *Bulletin of the Belgian Mathematical Society*, 6(1):147–154, 1999. Zentralblatt Math identifier: 0915.68143, Mathematical Reviews number (MathSciNet): MR1674709. Online available at <http://citeseer.ist.psu.edu/14550.html> and <http://projecteuclid.org/euclid.bbms/1103149975> [accessed 2007-11-05].
- [191] Niko Beerenwinkel, Lior Pachter, and Bernd Sturmfels. Epistasis and Shapes of Fitness Landscapes, 2006. Eprint arXiv:q-bio/0603034. Online available at <http://arxiv.org/abs/q-bio.PE/0603034> [accessed 2007-08-05].
- [192] Warren Sarle. What is overfitting and how can I avoid it? *Usenet FAQs: comp.ai.neural-nets FAQ*, 3: Generalization(3), May 13, 2007. Online available at <http://www.faqs.org/faqs/ai-faq/neural-nets/part3/section-3.html> [accessed 2008-02-29].
- [193] Murray Smith. *Neural Networks for Statistical Modeling*. John Wiley & Sons, Inc., New York, NY, USA / International Thomson Computer Press, Boston, USA, April 1993/1996. ISBN 0442013108, 1850328420, 978-1850328421.
- [194] Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1):1–58, 1992. ISSN 0899-7667.
- [195] Tom Dietterich from the Department of Computer Science, Oregon State University, Corvallis, OR 97331, USA. Overfitting and undercomputing in machine learning. *ACM Computing Surveys (CSUR)*, 27(3):326–327, 1995. ISSN 0360-0300. doi:10.1145/212094.212114. Online available at <http://doi.acm.org/10.1145/212094.212114> and <http://citeseer.ist.psu.edu/dietterich95overfitting.html> [accessed 2007-09-13].
- [196] Igor V. Tetko, David J. Livingstone, and Alexander I. Luik. Neural network studies, 1. Comparison of overfitting and overtraining. *Journal of Chemical Information and Computer Sciences*, 35(5):826–833, 1995.
- [197] Charles X. Ling from the Department of Computer Science; University of Western Ontario. Overfitting and generalization in learning discrete patterns. *Neurocomputing*, 8(3):341–347, August 1995. Optimization and Combinatorics, Part I-III, Online available at [http://dx.doi.org/10.1016/0925-2312\(95\)00050-G](http://dx.doi.org/10.1016/0925-2312(95)00050-G) and <http://citeseer.ist.psu.edu/514876.html> [accessed 2007-09-13].
- [198] Steve Lawrence and C. Lee Giles from the NEC Research Institute. Overfitting and Neural Networks: Conjugate Gradient and Backpropagation. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks (IJCNN'00)*, volume 1, pages 1114–1119. IEEE Computer Society, July 24–27, 2000, Como, Italy. Online available at <http://citeseer.ist.psu.edu/lawrence00overfitting.html> [accessed 2007-09-13].

- [199] Warren S. Sarle from the SAS Institute Inc., SAS Campus Drive, Cary, NC 27531, USA. Stopped Training and Other Remedies for Overfitting. In *Proceedings of the 27th Symposium on the Interface: Computing Science and Statistics*, pages 352–360, June 23, 1995, Pittsburgh, Pennsylvania. Online available at <http://citeseer.ist.psu.edu/71530.html> and <ftp://ftp.sas.com/pub/neural/inter95.ps.Z> [accessed 2007-09-13].
- [200] Paul L. Rosin and Freddy Fierens from the Institute for Remote Sensing Applications; Joint Research Centre, I-21020 Ispra (VA), Italy. Improving Neural Network Generalisation. In *Proceedings of the International Geoscience and Remote Sensing Symposium, IGARSS'95*, volume 1, July 10–14, 1995, Florenz, Italy. Online available at <http://citeseer.ist.psu.edu/165458.html> and <http://users.cs.cf.ac.uk/Paul.Rosin/resources/papers/overfitting.pdf> [accessed 2007-09-13].
- [201] Michael J. Kearns, Yishay Mansour, Andrew Y. Ng, and Dana Ron. An Experimental and Theoretical Comparison of Model Selection Methods. In *COLT'95: Proceedings of the eighth annual conference on Computational learning theory*, pages 21–30. ACM Press, New York, NY, USA, 1995, Santa Cruz, California, United States. ISBN 0-89791-723-5. doi:10.1145/225298.225301. Online available at <http://www.cis.upenn.edu/~mkearns/papers/ms.pdf> and <http://doi.acm.org/10.1145/225298.225301> [accessed 2008-03-02]. See also [1917].
- [202] Jürgen Branke from the Institute AIFB University of Karlsruhe, D-76128 Karlsruhe, Germany. Creating Robust Solutions by Means of Evolutionary Algorithms. In *Parallel Problem Solving from Nature – PPSN V*, 1998. In proceedings [588].
- [203] Shigeyoshi Tsutsui, Ashish Ghosh, and Yoshiji Fujimoto. A Robust Solution Searching Scheme in Genetic Search. In *PPSN IV: Proceedings of the 4th International Conference on Parallel Problem Solving from Nature*, pages 543–552, 1996. In proceedings [589]. Online available at <http://www.hannan-u.ac.jp/~tsutsui/ps/ppsn-iv.pdf> [accessed 2008-07-19].
- [204] Dirk Wiesmann, Ulrich Hammel, and Thomas Bäck. Robust design of multilayer optical coatings by means of evolutionary strategies. Sonderforschungsbereich (sfb) 531, Universität Dortmund, March 31 1998. Online available at <http://hdl.handle.net/2003/5348> and <http://citeseer.ist.psu.edu/325801.html> [accessed 2008-07-19]. See also [205].
- [205] Dirk Wiesmann, Ulrich Hammel, and Thomas Bäck. Robust design of multilayer optical coatings by means of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 2(4):162–167, November 1998. ISSN 1089-778X. doi:10.1109/4235.738986. CODEN: ITEVF5, INSPEC Accession Number: 6149966. See also [204].
- [206] Hisashi Handa, Dan Lin, Lee Chapman, and Xin Yao. Robust Solution of Salting Route Optimisation Using Evolutionary Algorithms. In *Evolutionary Computation, proceedings of CEC 2006*, pages 3098–3105, 2006. doi:10.1109/CEC.2006.1688701. Online available at http://escholarship.lib.okayama-u.ac.jp/industrial_engineering/2/ [accessed 2008-06-19].
- [207] Shigeyoshi Tsutsui and Ashish Ghosh from the Dept. of Manage. & Inf. Sci., Hannan Univ., Osaka. Genetic algorithms with a robust solution searching scheme. *IEEE Transactions on Evolutionary Computation*, 1(3):201–208, September 1997. ISSN 1089-778X. doi:10.1109/4235.661550. CODEN: ITEVF5, INSPEC Accession Number: 5858813.
- [208] Genichi Taguchi. *Introduction to Quality Engineering: Designing Quality into Products and Processes*. Asian Productivity Organization / American Supplier Institute Inc. / Quality Resources / Productivity Press Inc., January 1986. ISBN 9283310837, 978-9283310839 / 9283310845. Translation of Sekkeisha no tame no hinshitsu kanri.
- [209] Horst Greiner from the Philips Forschungslaboratorien GmbH, P.O. Box 1980, Weißhausstraße, Aachen D-52066, Germany. Robust Filter Design by Stochastic

- Optimization. *Proceedings of SPIE (The International Society for Optical Engineering)*, 2253(150):150–161, November 1994. doi:10.1117/12.192107. Advances in Design and Design Techniques: Posters, June 5, 1994, Grenoble, France, Optical Interference Coatings, Florin Abeles, ed.
- [210] Horst Greiner from the Philips Forschungslaboratorien GmbH, P.O. Box 1980, Weißhausstraße, Aachen D-52066, Germany. Robust optical coating design with evolutionary strategies. *Applied Optics*, 35(28):5477–5483, October 1, 1996.
- [211] J. Michael Fitzpatrick and John J. Grefenstette from the Computer Science Department, Vanderbilt University, 37235 Nashville, Tennessee, USA. Genetic Algorithms in Noisy Environments. *Machine Learning*, 3(2–3):101–120, October 1988. ISSN 0885-6125 (Print) 1573-0565 (Online). doi:10.1007/BF00113893. Online available at <http://www.springerlink.com/content/n6w328441t63q374/fulltext.pdf> [accessed 2008-07-19].
- [212] Yasuhito Sano and Hajime Kita. Optimization of Noisy Fitness Functions by Means of Genetic Algorithms Using History of Search. In *PPSN VI: Proceedings of the 6th International Conference on Parallel Problem Solving from Nature*, pages 571–580, 2000. doi:10.1007/3-540-45356-3_56. In proceedings [587].
- [213] Yasuhito Sano and Hajime Kita. Optimization of Noisy Fitness Functions by means of Genetic Algorithms using History of Search with Test of Estimation. In *Congress on Evolutionary Computation*, pages 360–365, 2002. In proceedings [499].
- [214] Hajime Kita and Yasuhito Sano. Genetic Algorithms for Optimization of Noisy Fitness Functions and Adaptation to Changing Environments. In *2003 Joint Workshop of Hayashibara Foundation and 2003 Workshop on Statistical Mechanical Approach to Probabilistic Information Processing (SMAPIP)*, July 14–15, 2003, Okayama, Japan. Online available at <http://www.smapip.is.tohoku.ac.jp/~smapip/2003/hayashibara/proceedings/HajimeKita.pdf> [accessed 2008-07-19].
- [215] Hans-Georg Beyer. Toward a Theory of Evolution Strategies: Some Asymptotical Results from the $(1, +\lambda)$ -Theory. *Evolutionary Computation*, 1(2 (Summer)):165–188, 1993. ISSN 1063-6560, 1530-9304. Online available at <http://citeseer.ist.psu.edu/beyer93towards.html> [accessed <http://ls11-www.cs.uni-dortmund.de/~beyer/coll/Bey93a/Bey93a.ps>] 2008-07-19. See also [1918].
- [216] Thomas Bäck and Ulrich Hammel from the University of Dortmund, Department of Computer Science, LS XI, D-44221, Dortmund, Germany. Evolution Strategies Applied to Perturbed Objective Functions. In *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence*, volume 1, pages 40–45, 1994. doi:10.1109/ICEC.1994.350045. In proceedings [507]. Online available at <http://citeseer.ist.psu.edu/16973.html> [accessed 2008-07-19].
- [217] Ulrich Hammel and Thomas Bäck. Evolution Strategies on Noisy Functions: How to Improve Convergence Properties. In *Parallel Problem Solving from Nature – PPSN III, International Conference on Evolutionary Computation. Proceedings of the Third Conference on Parallel Problem Solving from Nature*, pages 159–168, 1994. In proceedings [590]. Online available at <http://citeseer.ist.psu.edu/24308.html> [accessed 2008-07-19].
- [218] Hui Pan, Ling Wang, and Bo Liu from the Department of Automation, Tsinghua University, Beijing 100084, China. Particle swarm optimization for function optimization in noisy environment. *Applied Mathematics and Computation*, 181(2):908–919, October 15, 2006. ISSN 0096-3003. doi:10.1016/j.amc.2006.01.066. CODEN AMHCBQ.
- [219] Lin Han and Xingshi He. A Novel Opposition-Based Particle Swarm Optimization for Noisy Problems. In *ICNC'07: Proceedings of the Third International Conference on Natural Computation*, volume 3, pages 624–629, 2007. doi:10.1109/ICNC.2007.119. In proceedings [312].
- [220] Brad L. Miller and David E. Goldberg from the University of Illinois at Urbana/Champaign Urbana, IL 61801. Genetic algorithms, selection schemes, and the

- varying effects of noise. *Evolutionary Computation*, 4(2):113–131, Summer 1996. ISSN 1063-6560. doi:10.1162/evco.1996.4.2.113. See also [221]. Online available at <http://citeseer.ist.psu.edu/110556.html> and <http://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=4225507A21F8083F025E28499309D1FB?doi=10.1.1.31.3449&rep=rep1&type=pdf> [accessed 2008-07-21].
- [221] Brad L. Miller and David E. Goldberg. Genetic Algorithms, Tournament Selection, and the Effects of Noise. IlliGAL Report 95006, Illinois Genetic Algorithms Laboratory, Department of General Engineering, University of Illinois, Urbana-Champaign, 117 Transportation Building, 104 South Mathews Avenue, Urbana, IL 61801, USA, July 1995. See also [220]. Online available at <http://citeseer.ist.psu.edu/86198.html> [accessed <http://www.illgal.uiuc.edu/pub/papers/IlliGALs/95006.ps.z>] 2007-07-28.
- [222] Jack Y. B. Lee and P. C. Wong from the Advanced Network Systems Laboratory Department of Information Engineering, The Chinese University of Hong Kong, Hongkong. The effect of function noise on GP efficiency. In *Progress in Evolutionary Computation*, pages 1–16, 1995. doi:10.1007/3-540-60154-6_43. In proceedings [1919].
- [223] L. S. Gurin and Leonard A. Rastrigin. Convergence of the random search method in the presence of noise. *Automation and Remote Control*, 26:1505–1511, 1965.
- [224] Jürgen Branke, Erdem Salihoglu, and Şima Uyar. Towards an analysis of dynamic environments. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1433–1440, 2005. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068237> and <http://www.citeulike.org/user/denizinho/article/995596> [accessed 2007-08-19].
- [225] Ronald Walter Morrison and Kenneth Alan De Jong. A test problem generator for non-stationary environments. In *Proceedings of the 1999 Congress on Evolutionary Computation, CEC 99*, volume 3, pages 2047–2053, 1999. doi:10.1109/CEC.1999.785526. In proceedings [502].
- [226] Hendrik Richter from the HTWK Leipzig, Fachbereich Elektrotechnik und Informationstechnik, Institut Mess-, Steuerungs- und Regelungstechnik, Postfach 30 11 66, D-04125 Leipzig, Germany. Behavior of Evolutionary Algorithms in Chaotically Changing Fitness Landscapes. In *Proceedings of 8th International Conference on Parallel Problem Solving from Nature, PPSN VIII*, pages 111–120, 2004. In proceedings [585].
- [227] Jürgen Branke. The Moving Peaks Benchmark. Technical Report, Institute AIFB, University of Karlsruhe, D-76128 Karlsruhe, Germany, 1999. Online available at <http://www.aifb.uni-karlsruhe.de/~jbr/MovPeaks/> [accessed 2007-08-19].
- [228] Shengxiang Yang, Yew-Soon Ong, and Yaochu Jin, editors. *Evolutionary Computation in Dynamic and Uncertain Environments*, volume 51(XXIII) of *Studies in Computational Intelligence*. Springer, 2007. ISBN 978-3-540-49772-1. Presentation online available at http://www.soft-computing.de/Jin_CEC04T.pdf.gz [accessed 2007-08-19].
- [229] Krzysztof Trojanowski. *Evolutionary Algorithms with Redundant Genetic Material for Non-Stationary Environments*. PhD thesis, Instytut Podstaw Informatyki PAN, Institute of Computer Science, Warsaw, University of Technology, Poland, 1994. Advisor: Zbigniew Michalewicz.
- [230] Claus O. Wilke. *Evolutionary Dynamics in Time-Dependent Environments*. PhD thesis, Fakultät für Physik und Astronomie, Ruhr-Universität Bochum, Shaker Verlag, Aachen, July 1999. Online available at <http://wlab.biosci.utexas.edu/~wilke/ps/PhD.ps.gz> [accessed 2007-08-19].
- [231] Jürgen Branke. *Evolutionary Optimization in Dynamic Environments*. PhD thesis, Universität Karlsruhe (TH), Fakultät für Wirtschaftswissenschaften, Universität Karlsruhe (TH), Institut AIFB, D-76128 Karlsruhe, 2000. Advisors: H. Schmeck and G. Bol and Lothar Thiele. See also [232].

- [232] Jürgen Branke. *Evolutionary Optimization in Dynamic Environments*, volume 3 of *Genetic Algorithms and Evolutionary Computation*. Kluwer Academic Publishers, December 2001. ISBN 978-0792376316. See also [231].
- [233] Ronald Walter Morrison. *Designing Evolutionary Algorithms for Dynamic Environments*. PhD thesis, George Mason University, USA, 2002. Advisor: Kenneth Alan De Jong. See also [234].
- [234] Ronald Walter Morrison. *Designing Evolutionary Algorithms for Dynamic Environments*, volume 24(1) of *Natural Computing*. Springer, Berlin, June 2004. ISBN 978-3540212317. See also [233].
- [235] Victoria S. Aragón and Susana C. Esquivel. An Evolutionary Algorithm to Track Changes of Optimum Value Locations in Dynamic Environments. *Journal of Computer Science & Technology*, 4(3)(12), October 2004. Online available at <http://journal.info.unlp.edu.ar/Journal/journal12/papers.html> [accessed 2007-08-19].
- [236] Helen G. Gobb and John J. Grefenstette. Genetic Algorithms for Tracking Changing Environments. In *Proceedings of 5th ICGA*, pages 523–529, 1993. In proceedings [748]. Online available at <ftp://ftp.aic.nrl.navy.mil/pub/papers/1993/AIC-93-004.ps> [accessed 2008-07-19].
- [237] Naoki Mori, Hajime Kita, and Yoshikazu Nishikawa from the Department of Electrical Engineering, Graduate School of Engineering, Kyoto University, Yoshidahonmachi, Sakyo, 606-01 Kyoto, Japan. Adaptation to a changing environment by means of the thermodynamical genetic algorithm. In *Parallel Problem Solving from Nature – PPSN IV*, pages 513–522, 1996. doi:10.1007/3-540-61723-X_1015. In proceedings [589].
- [238] Naoki Mori, Seiji Imanishi, Hajime Kita, and Yoshikazu Nishikawa. Adaptation to Changing Environments by Means of the Memory Based Thermodynamical Genetic Algorithm. In *ICGA*, pages 299–306, 1997. In proceedings [746].
- [239] Naoki Mori, Hajime Kita, and Yoshikazu Nishikawa. Adaptation to a changing environment by means of the feedback thermodynamical genetic algorithm. In *Parallel Problem Solving from Nature – PPSN V*, pages 149–158, 1998. doi:10.1007/BFb0056858. In proceedings [588].
- [240] Anthony Jack Carlisle. *Applying the Particle Swarm Optimizer to Non-Stationary Environments*. PhD thesis, Graduate Faculty of Auburn University, December 2002. Advisors: Gerry V. Dozier. Online available at <http://antho.huntingdon.edu/publications/default.html> [accessed 2007-08-19].
- [241] Anthony Jack Carlisle and Gerry V. Dozier. Tracking changing extrema with adaptive particle swarm optimizer. In *Proceedings of the 5th Biannual World Automation Congress, WAC 2002*, volume 13, pages 265–270, June 9–13, 2002, Orlando, Florida, USA. doi:10.1109/WAC.2002.1049555. Online available at <http://antho.huntingdon.edu/publications/default.html> [accessed 2007-08-19].
- [242] Xiaodong Li, Jürgen Branke, and Tim Blackwell. Particle swarm with speciation and adaptation in a dynamic environment. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 51–58, 2006. doi:10.1145/1143997.1144005. In proceedings [553]. Online available at <http://doi.acm.org/10.1145/1143997.1144005> [accessed 2007-08-19].
- [243] Guanyu Pan, Quansheng Dou, and Xiaohua Liu. Performance of two Improved Particle Swarm Optimization In Dynamic Optimization Environments. In *ISDA'06: Proceedings of the Sixth International Conference on Intelligent Systems Design and Applications (ISDA'06)*, volume 2, pages 1024–1028, 2006, Jinan, China. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2528-8.
- [244] Tim Blackwell. Particle Swarm Optimization in Dynamic Environments. In *Evolutionary Computation in Dynamic and Uncertain Environments*, chapter 2, pages 29–52. Springer, 2007. In collection [228]. See also [228]. Online available at <http://igor.gold.ac.uk/~mas01tb/papers/PSOdynenv.pdf> [accessed 2007-08-19].

- [245] Rui Mendes and Arvind S. Mohais. DynDE: a differential evolution for dynamic optimization problems. In *Proceedings of 2005 IEEE Congress on Evolutionary Computation*, volume 3, pages 2808–2815, 2005. In proceedings [496]. Online available at <http://www.di.uminho.pt/~rcm/publications/DynDE.pdf> and <http://en.scientificcommons.org/8412355> [accessed 2007-08-13].
- [246] Nelson Wu. Differential Evolution for optimisation in dynamic environments. Technical Report, School of Computer Science and Information Technology, RMIT University, November 2006. Online available at <http://yallara.cs.rmit.edu.au/~newu/portfolio.html> [accessed 2007-08-19].
- [247] Michael Gunttsch and Martin Middendorf. Pheromone Modification Strategies for Ant Algorithms Applied to Dynamic TSP. In *Proceedings of the EvoWorkshops on Applications of Evolutionary Computing*, pages 213–222, 2001. In proceedings [535]. Online available at <http://pacosy.informatik.uni-leipzig.de/pv/Personen/middendorf/Papers/> [accessed 2007-08-19].
- [248] Michael Gunttsch, Martin Middendorf, and Hartmut Schmeck. An Ant Colony Optimization Approach to Dynamic TSP. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 860–867, 2001. In proceedings [568]. Online available at <http://citeseer.ist.psu.edu/693476.html> [accessed 2007-08-19].
- [249] Nicolas J. Radcliffe from the Edinburgh Parallel Computing Centre, University of Edinburgh, Kings Buildings, EH9 3JZ, Scotland. The Algebra of Genetic Algorithms. *Annals of Mathematics and Artificial Intelligence*, 10(4), December 1994. ISSN 1012-2443 (print version), 1573-7470 (electronic version). doi:10.1007/BF01531276. Online available at <http://www.springerlink.com/content/w20613t556788tr0/fulltext.pdf> [accessed 2008-08-10]. Also technical report TR92-11, 1992 from Edinburgh Parallel Computing Centre, University of Edinburgh, King’s Buildings, EH9 3JZ, Scotland. Online available at <http://citeseer.ist.psu.edu/radcliffe94algebra.html> and <http://users.breathes.com/njr/formaPapers.html> [accessed 2007-07-28].
- [250] Mario Köppen, David H. Wolpert, and William G. Macready. Remarks on a recent paper on the “no free lunch” theorems. *IEEE Transactions on Evolutionary Computation*, 5(3):295–296, June 2001. doi:10.1109/4235.930318. Online available at <http://www.no-free-lunch.org/KoWM01.pdf> and <http://ti.arc.nasa.gov/people/dhw/papers/76.ps> [accessed 2008-03-28].
- [251] Christian Igel and Marc Toussaint. On classes of functions for which No Free Lunch results hold. *Information Processing Letters*, 86(6):317–321, June 30, 2003. ISSN 0020-0190. doi:10.1016/S0020-0190(03)00222-9. Online available at [http://dx.doi.org/10.1016/S0020-0190\(03\)00222-9](http://dx.doi.org/10.1016/S0020-0190(03)00222-9) and <http://citeseer.ist.psu.edu/623026.html> [accessed 2008-03-28].
- [252] Christian Igel and Marc Toussaint. Recent Results on No-Free-Lunch Theorems for Optimization, 2003. Online available at <http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0303032> [accessed 2008-03-28].
- [253] John Henry Holland. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor, 1975. ISBN 978-026258111. Reprinted by MIT Press, April 1992.
- [254] Kenneth Alan De Jong. *An analysis of the behavior of a class of genetic adaptive systems*. PhD thesis, Computer and Communication Sciences, University of Michigan, August 1975. Order number AAI7609381. Doctoral Committee: John Henry Holland, Larry K. Flanigan, Richard A. Volz, Bernhard P. Zeigler. Online available at http://cs.gmu.edu/~eclab/kdj_thesis.html [accessed 2007-08-17].
- [255] Karsten Weicker. *Evolutionäre Algorithmen*. Leitfäden der Informatik. B. G. Teubner GmbH, Stuttgart/Leipzig/Wiesbaden, March 2002. ISBN 3-519-00362-7.
- [256] Patrick David Surry. *A Prescriptive Formalism for Constructing Domain-specific Evolutionary Algorithms*. PhD thesis, University of Edinburgh, 1998. Online available

- at <http://citeseer.ist.psu.edu/258999.html> [accessed 2008-07-22].
- [257] Nicholas J. Radcliffe. Non-Linear Genetic Representations. In *Parallel problem solving from nature 2*, pages 259–268, 1992. In proceedings [591]. Online available at <http://citeseer.ist.psu.edu/radcliffe92nonlinear.html> and <http://users.breathe.com/njr/papers/ppsn92.pdf> [accessed 2007-09-05].
 - [258] Nicholas J. Radcliffe from the Department of Physics, University of Edinburgh, King's Buildings, EH9 3JZ, Scotland. Equivalence Class Analysis of Genetic Algorithms. *Complex Systems*, 5(2):183–205, 1991. Online available at <http://citeseer.ist.psu.edu/608745.html> [accessed 2007-07-28].
 - [259] Michael D. Vose from the University of Tennessee, Department of Computer Science, 107 Ayres Hall, Knoxville, TN 37996, USA. Generalizing the notion of schema in genetic algorithms. *Artificial Intelligence*, 50(3):385–396, 1991. ISSN 0004-3702. doi:10.1016/0004-3702(91)90019-G.
 - [260] Xuan Hoai Nguyen. *A Flexible Representation for Genetic Programming: Lessons from Natural Language Processing*. PhD thesis, School of Information Technology and Electrical Engineering University College, University of New South Wales, Australian Defence Force Academy, December 2004. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/hoai_thesis.html and http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/hoai_thesis.tar.gz [accessed 2007-07-15].
 - [261] Charles Campbell Palmer and Aaron Kershenbaum. Representing trees in genetic algorithms. In *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence*, volume 1, pages 379–384, 1994. In proceedings [507], see also [262].
 - [262] Charles Campbell Palmer and Aaron Kershenbaum. Representing trees in genetic algorithms. In *Handbook of Evolutionary Computation*. Oxford University Press, 1997. In collection [410], see also [261]. Online available at <http://citeseer.ist.psu.edu/palmer94representing.html> and <http://www.cs.cinvestav.mx/~constraint/> [accessed 2007-08-12].
 - [263] Charles Campbell Palmer. *An approach to a problem in network design using genetic algorithms*. PhD thesis, Polytechnic University, New York, NY, Polytechnic University, New York, NY, April 1994. Supervisors: Aaron Kershenbaum, Susan Flynn Hummel, Richard M. van Slyke, Robert R. Boorstyn. UMI Order No. GAX94-31740. Appears also as article in Wiley InterScience, Networks, Volume 26, Issue 3, Pages 151–163, Online available at <http://citeseer.ist.psu.edu/palmer95approach.html> [accessed 2007-08-12].
 - [264] Simon Ronald from the Program & Data Optimisation Group, Adelaide University, SA, Australia. Robust Encodings in Genetic Algorithms: A Survey of Encoding Issues. In *IEEE Forth International Conference on Evolutionary Computation (IEEE/ICEC'97)*, pages 43–48, 1997. doi:10.1109/ICEC.1997.592265. In proceedings [504].
 - [265] V. Bhaskar, Santosh K. Gupta, and A.K. Ray. Applications of multiobjective optimization in chemical engineering. *Reviews in Chemical Engineering*, 16(1):1–54, 2000.
 - [266] Nikolaos V. Sahinidis and Mohit Tawarmalani. Applications of Global Optimization to Process and Molecular Design. *Computers and Chemical Engineering*, 24(9-10): 2157–2169, October 1, 2000. Online available at <http://citeseer.ist.psu.edu/397733.html> [accessed 2007-09-01].
 - [267] Christodoulos A. Floudas and Panos M. Pardalos, editors. *Frontiers in Global Optimization*. Nonconvex Optimization and Its Applications. Springer US, November 30, 2003. ISBN 978-1402076992.
 - [268] Christodoulos A. Floudas, editor. *Deterministic Global Optimization: Theory, Methods and Applications*, volume 37 of *Nonconvex Optimization and Its Applications*.

- Springer-Verlag GmbH, 1999. ISBN 978-0-7923-6014-8. Partly online available at <http://books.google.de/books?id=qZSpq27Ts0cC> [accessed 2008-03-10].
- [269] Lothar Birk, Günther F. Clauss, and June Y. Lee from the Institute of Land and Sea Transport Systems, Technical University of Berlin, Berlin, Germany. Practical Application of Global Optimization to the Design of Offshore Structures. In *Proceedings of OMAE04, 23rd International Conference on Offshore Mechanics and Arctic Engineering*, June 20–25, 2004, Vancouver, British Columbia, Canada. OM AE2004-51225. Online available at <http://130.149.35.79/downloads/publikationen/2004/OMAE2004-51225.pdf> [accessed 2007-09-01].
- [270] G. Dzemyda, V. Saltenis, and A. Zilinskas, editors. *Stochastic and Global Optimization*. Nonconvex Optimization and Its Applications. Springer-Verlag GmbH, March 1, 2002. ISBN 978-1402004841.
- [271] Pravir K. Chawdry, Rajkumar Roy, and Raj K. Pant, editors. *Soft Computing in Engineering Design and Manufacturing*. Springer-Verlag, August 1998. ISBN 978-3540762140. Partly online available at <http://books.google.de/books?id=mxCP1mSj0lSc> [accessed 2008-06-07].
- [272] Max Jerrell from the Northern Arizona University. Applications of Public Global Optimization Software to Difficult Econometric Functions. In *Computing in Economics and Finance 200*. Society for Computational Economics, 2000, Departament d'Economia i Empresa, Universitat Pompeu Fabra, Ramon Trias Fargas, 25,27, 08005, Barcelona, Spain. Number 161. Online available at <http://econpapers.repec.org/paper/scsescf0/161.htm> [accessed 2007-09-01].
- [273] E. A. Galperin and E. J. Kansa. Application of global optimization and radial basis functions to numerical solutions of weakly singular volterra integral equations. *Computers & Mathematics with Applications*, 43(3-5):491–499, February–March 2002. doi:10.1016/S0898-1221(01)00300-5. Online available at [http://dx.doi.org/10.1016/S0898-1221\(01\)00300-5](http://dx.doi.org/10.1016/S0898-1221(01)00300-5) [accessed 2007-09-01].
- [274] Saswatee Banerjee and Lakshminarayan N. Hazra from the University of Calcutta. Thin lens design of Cooke triplet lenses: application of a global optimization technique. In Kevin D. Bell, Michael K. Powers, and Jose M. Sasian, editors, *Proceedings of SPIE, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference – Novel Optical Systems and Large-Aperture Imaging*, volume 3430, pages 175–183, December 1998. doi:10.1117/12.332474.
- [275] Chao-Hsi Tsao and Jyh-Long Chern from the National Chiao Tung University, Institute of Electro-Optical Engineering, Microelectronics and Information System, Research Center, Department of Photonics, Hsinchu 300, Taiwan. Application of a global optimization process to the design of pickup heads for compact and digital versatile disks. *Optical Engineering*, 45(10):103001, October 2006. doi:10.1117/1.2361281.
- [276] Yacov Y. Haimes, Warren Hall, and Herbert T. Freedman. *Multiobjective Optimization in Water Resource Systems*. Elsevier, New York, 1975. ASIN: B000UUMGXE.
- [277] David W. Corne, Martin J. Oates, and George D. Smith, editors. *Telecommunications Optimization: Heuristic and Adaptive Techniques*. John Wiley & Sons, Ltd., September 2000. ISBN 0471988553, 9780471988557, 9780470841631. doi:10.1002/047084163X.
- [278] *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, AAAI*, July 22–26, 2007, Vancouver, British Columbia, Canada. AAAI Press. ISBN 978-1-57735-323-2. See <http://www.aaai.org/Conferences/AAAI/aaai07.php> [accessed 2007-09-06].
- [279] *Proceedings, The Twenty-First National Conference on Artificial Intelligence (AAAI) and the Eighteenth Innovative Applications of Artificial Intelligence Conference (IAAI)*, July 16–20, 2006, Boston, Massachusetts, USA. AAAI Press. See <http://www.aaai.org/Conferences/AAAI/aaai06.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai06.php> [accessed 2007-09-06].

- [280] Subbarao Kambhampati Manuela M. Veloso, editor. *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI) and the Seventeenth Innovative Applications of Artificial Intelligence Conference (IAAI)*, July 9–13, 2005, Pittsburgh, Pennsylvania, USA. AAAI Press/The MIT Press. ISBN 1-57735-236-X. See <http://www.aaai.org/Conferences/AAAI/aaai05.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai06.php> [accessed 2007-09-06].
- [281] Deborah L. McGuinness and George Ferguson, editors. *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI) and the Sixteenth Conference on Innovative Applications of Artificial Intelligence (IAAI)*, July 25–29, 2004, San Jose, California, USA. AAAI Press/The MIT Press. ISBN 0-262-51183-5. See <http://www.aaai.org/Conferences/AAAI/aaai04.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai04.php> [accessed 2007-09-06].
- [282] *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI) and Fourteenth Conference on Innovative Applications of Artificial Intelligence (IAAI)*, July 28–August 1, 2002, Edmonton, Alberta, Canada. AAAI Press. See <http://www.aaai.org/Conferences/AAAI/aaai02.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai02.php> [accessed 2007-09-06].
- [283] *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI) and Twelfth Conference on Innovative Applications of Artificial Intelligence (IAAI)*, July 30–August 3, 2000, Austin, Texas, USA. AAAI Press/The MIT Press. ISBN 0-262-51112-6. See <http://www.aaai.org/Conferences/AAAI/aaai00.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai00.php> [accessed 2007-09-06].
- [284] *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI) and Eleventh Conference on Innovative Applications of Artificial Intelligence (IAAI)*, July 18–22, 1999, Orlando, Florida, USA. AAAI Press/The MIT Press. ISBN 0-262-51106-1. See <http://www.aaai.org/Conferences/AAAI/aaai99.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai99.php> [accessed 2007-09-06].
- [285] *Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI 98, IAAI 98*, July 26–30, 1998, Madison, Wisconsin, USA. AAAI Press/The MIT Press. ISBN 0-262-51098-7. See <http://www.aaai.org/Conferences/AAAI/aaai98.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai98.php> [accessed 2007-09-06].
- [286] *Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Innovative Applications of Artificial Intelligence Conference, AAAI 97, IAAI 97*, July 27–31, 1997, Providence, Rhode Island, USA. AAAI Press/The MIT Press. ISBN 0-262-51095-2. See <http://www.aaai.org/Conferences/AAAI/aaai97.php> [accessed 2007-09-06] and <http://www.aaai.org/Conferences/IAAI/iaai97.php> [accessed 2007-09-06].
- [287] *Deep Blue Versus Kasparov: The Significance for Artificial Intelligence, Collected Papers from the 1997 AAAI Workshop*, 1997. AAAI Press. ISBN 1-57735-031-6.
- [288] *Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference, AAAI 96, IAAI 96*, August 4–8, 1996, Portland, Oregon, USA. AAAI Press/The MIT Press. ISBN 0-262-51091-X. 2 volumes, see <http://www.aaai.org/Conferences/AAAI/aaai96.php> [accessed 2007-09-06], [289], and <http://www.aaai.org/Conferences/IAAI/iaai96.php> [accessed 2007-09-06].
- [289] *Agent Modeling, Papers from the 1996 AAAI Workshop*, 1996. AAAI Press. ISBN 1-57735-008-1. See [288].
- [290] *Proceedings of the 12th National Conference on Artificial Intelligence, AAAI*, July 31–August 4, 1994, Convention Center, Seattle, WA, USA. AAAI Press. ISBN 0-262-61102-3. 2 volumes, see <http://www.aaai.org/Conferences/AAAI/aaai94.php> [accessed 2007-09-06].

- [291] *Proceedings of the 11th National Conference on Artificial Intelligence, AAAI*, July 11–15, 1993, Washington, DC, USA. The AAAI Press/The MIT Press. ISBN 0-262-51071-5. See <http://www.aaai.org/Conferences/AAAI/aaai93.php> [accessed 2007-09-06].
- [292] William R. Swartout, editor. *Proceedings of the 10th National Conference on Artificial Intelligence, AAAI*, July 12–16, 1992, San Jose, California, USA. The AAAI Press/The MIT Press. ISBN 0-262-51063-4. See <http://www.aaai.org/Conferences/AAAI/aaai92.php> [accessed 2007-09-06].
- [293] *Proceedings of the 9th National Conference on Artificial Intelligence, AAAI*, July 14–19, 1991, Anaheim, California, USA. AAAI Press/The MIT Press. ISBN 0-262-51059-6. 2 volumes, see <http://www.aaai.org/Conferences/AAAI/aaai91.php> [accessed 2007-09-06].
- [294] *Proceedings of the 8th National Conference on Artificial Intelligence, AAAI*, July 29–August 3, 1990, Boston, Massachusetts, USA. AAAI Press/The MIT Press. ISBN 0-262-51057-X. 2 volumes, see <http://www.aaai.org/Conferences/AAAI/aaai90.php> [accessed 2007-09-06].
- [295] *Proceedings of the 7th National Conference on Artificial Intelligence, AAAI*, August 21–26, 1988, St. Paul, Minnesota, USA. AAAI Press/The MIT Press. ISBN 0-262-51055-3. See <http://www.aaai.org/Conferences/AAAI/aaai88.php> [accessed 2007-09-06].
- [296] K. S. H. Forbus, editor. *Proceedings of the 6th National Conference on Artificial Intelligence, AAAI*, July 1987, Seattle, WA, USA. Morgan Kaufmann, San Francisco, CA, USA. See <http://www.aaai.org/Conferences/AAAI/aaai87.php> [accessed 2007-09-06].
- [297] *Proceedings of the 5th National Conference on Artificial Intelligence, AAAI*, volume Science (volume 1), August 11–15, 1986, Philadelphia, PA, USA. Morgan Kaufmann, San Francisco, CA, USA. See <http://www.aaai.org/Conferences/AAAI/aaai86.php> [accessed 2007-09-06] and also [298].
- [298] *Proceedings of the 5th National Conference on Artificial Intelligence, AAAI*, volume Engineering (volume 2), August 11–15, 1986, Philadelphia, PA, USA. Morgan Kaufmann, San Francisco, CA, USA. See <http://www.aaai.org/Conferences/AAAI/aaai86.php> [accessed 2007-09-06] and also [297].
- [299] Ronald J. Brachman, editor. *Proceedings of the National Conference on Artificial Intelligence, AAAI*, August 6–10, 1984, Austin, TX, USA. AAAI Press. ISBN 0-262-51053-7. See <http://www.aaai.org/Conferences/AAAI/aaai84.php> [accessed 2007-09-06].
- [300] Michael R. Genesereth, editor. *Proceedings of the National Conference on Artificial Intelligence, AAAI*, August 22–26, 1983, Washington, DC, USA. AAAI Press. ISBN 0-262-51052-9. See <http://www.aaai.org/Conferences/AAAI/aaai83.php> [accessed 2007-09-06].
- [301] David L. Waltz, editor. *Proceedings of the National Conference on Artificial Intelligence, AAAI*, August 18–20, 1982, Pittsburgh, PA, USA. AAAI Press. ISBN 0-262-51051-2. See <http://www.aaai.org/Conferences/AAAI/aaai82.php> [accessed 2007-09-06].
- [302] Robert Balzer, editor. *Proceedings of the 1st Annual National Conference on Artificial Intelligence, AAAI*, August 18–21, 1980, Stanford University, California, USA. AAAI Press/MIT Press. ISBN 0-262-51050-2. See <http://www.aaai.org/Conferences/AAAI/aaai80.php> [accessed 2007-09-06].
- [303] Andreas König, Mario Köppen, Ajith Abraham, Christian Igel, and Nikola Kasabov, editors. *7th International Conference on Hybrid Intelligent Systems (HIS 2007)*, September 17–19, 2007, Fraunhofer-Center, University of Kaiserslautern, Kaiserslautern, Germany. IEEE Computer Society. ISBN 0-7695-2946-1. Library

- of Congress Number 2007936727, Product Number E2946. see <http://his07.hybridsystem.com/> [accessed 2007-09-01].
- [304] *6th International Conference on Hybrid Intelligent Systems (HIS 2006)*, December 13–15, 2006, AUT Technology Park, Auckland, New Zealand. IEEE Computer Society. ISBN 0-7695-2662-4. see <http://his06.hybridsystem.com/> [accessed 2007-09-01].
 - [305] Nadia Nedjah, Luiza de Macedo Mourelle, Ajith Abraham, and Mario Köppen, editors. *5th International Conference on Hybrid Intelligent Systems (HIS 2005)*, November 6–9, 2005, Pontifical Catholic University of Rio de Janeiro: PUC-Rio, Rio de Janeiro, Brazil. IEEE Computer Society. ISBN 0-7695-2457-5. see <http://his05.hybridsystem.com/> [accessed 2007-09-01].
 - [306] *4th International Conference on Hybrid Intelligent Systems (HIS 2004)*, December 5–8, 2005, Kitakyushu, Japan. IEEE Computer Society. ISBN 0-7695-2291-2. INSPEC Accession Number: 8470883.
 - [307] Ajith Abraham, Mario Köppen, and Katrin Franke, editors. *Design and Application of Hybrid Intelligent Systems, HIS03, Proceedings of the Third International Conference on Hybrid Intelligent Systems*, volume 105 of *Frontiers in Artificial Intelligence and Applications*, December 14–17, 2003, Monash Conference Centre, Level 7, 30 Collins Street, Melbourne, Australia. IOS Press. ISBN 1-58603-394-8. see <http://his03.hybridsystem.com/> [accessed 2007-09-01].
 - [308] Ajith Abraham, Javier Ruiz del Solar, and Mario Köppen, editors. *Soft Computing Systems – Design, Management and Applications, HIS 2002*, volume 87 of *Frontiers in Artificial Intelligence and Applications*, December 1–4, 2002, Santiago, Chile. IOS Press. ISBN 1-58603-297-6.
 - [309] Ajith Abraham and Mario Köppen, editors. *Hybrid Information Systems, Proceedings of the First International Workshop on Hybrid Intelligent Systems*, Advances in Soft Computing, December 11–12, 2002, Adelaide University’s main campus, Adelaide, Australia. Physica-Verlag. ISBN 3-7908-1480-6. see <http://his01.hybridsystem.com/> [accessed 2007-09-01].
 - [310] *Proceedings of the Third International Conference on Advances in Natural Computation, ICNC 2007*, volume 1, August 24–27, 2007, Haikou, China. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2875-9, 978-0-7695-2875-5. doi:10.1109/ICNC.2007.763. See also [311, 312, 313, 314].
 - [311] *Proceedings of the Third International Conference on Advances in Natural Computation, ICNC 2007*, volume 2, August 24–27, 2007, Haikou, China. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2875-9, 978-0-7695-2875-5. doi:10.1109/ICNC.2007.763. See also [310, 312, 313, 314].
 - [312] *Proceedings of the Third International Conference on Advances in Natural Computation, ICNC 2007*, volume 3, August 24–27, 2007, Haikou, China. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2875-9, 978-0-7695-2875-5. doi:10.1109/ICNC.2007.766. See also [310, 311, 313, 314].
 - [313] *Proceedings of the Third International Conference on Advances in Natural Computation, ICNC 2007*, volume 4, August 24–27, 2007, Haikou, China. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2875-9, 978-0-7695-2875-5. doi:10.1109/ICNC.2007.765. See also [310, 311, 312, 314].
 - [314] *Proceedings of the Third International Conference on Advances in Natural Computation, ICNC 2007*, volume 5, August 24–27, 2007, Haikou, China. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2875-9, 978-0-7695-2875-5. doi:10.1109/ICNC.2007.762. See also [310, 311, 312, 313].
 - [315] Licheng Jiao, Lipo Wang, Xinbo Gao, Jing Liu, and Feng Wu, editors. *Proceedings of the Second International Conference on Advances in Natural Computation, ICNC 2006, Part I*, volume 4221 of *Lecture Notes in Computer Science (LNCS)*, September

- ber 24–28, 2006, Xi'an, China. Springer Berlin / Heidelberg. ISBN 3-540-45901-4, 978-3-540-45901-9. doi:10.1007/11881070. See also [316].
- [316] Licheng Jiao, Lipo Wang, Xinbo Gao, Jing Liu, and Feng Wu, editors. *Proceedings of the Second International Conference on Advances in Natural Computation, ICNC 2006, Part II*, volume 4222 of *Lecture Notes in Computer Science (LNCS)*, September 24–28, 2006, Xi'an, China. Springer Berlin / Heidelberg. ISBN 3-540-45907-3, 978-3-540-45907-1. doi:10.1007/11881223. See also [315].
- [317] Lipo Wang, Ke Chen, and Yew-Soon Ong, editors. *Proceedings of the First International Conference on Advances in Natural Computation, ICNC 2005, Part I*, volume 3610 of *Lecture Notes in Computer Science (LNCS)*, August 27–29, 2005, Changsha, China. Springer Berlin / Heidelberg. ISBN 3-540-28323-4, 978-3-540-28323-2. doi:10.1007/11539087. See also [318, 319].
- [318] Lipo Wang, Ke Chen, and Yew-Soon Ong, editors. *Proceedings of the First International Conference on Advances in Natural Computation, ICNC 2005, Part II*, volume 3611 of *Lecture Notes in Computer Science (LNCS)*, August 27–29, 2005, Changsha, China. Springer Berlin / Heidelberg. ISBN 3-540-28325-0, 978-3-540-28325-6. doi:10.1007/11539117. See also [317, 319].
- [319] Lipo Wang, Ke Chen, and Yew-Soon Ong, editors. *Proceedings of the First International Conference on Advances in Natural Computation, ICNC 2005, Part III*, volume 3612 of *Lecture Notes in Computer Science (LNCS)*, August 27–29, 2005, Changsha, China. Springer Berlin / Heidelberg. ISBN 3-540-28320-X, 978-3-540-28320-1. doi:10.1007/11539902. See also [317, 318].
- [320] John Riedl and Randall W. Hill Jr., editors. *Proceedings of the Fifteenth Conference on Innovative Applications of Artificial Intelligence*, August 12–14, 2003, Acapulco, México. AAAI. ISBN 1-57735-188-6. See <http://www.aaai.org/Conferences/IAAI/iaai03.php> [accessed 2007-09-06].
- [321] Haym Hirsh and Steve A. Chien, editors. *Proceedings of the Thirteenth Innovative Applications of Artificial Intelligence Conference*, August 7–9, 2001, Seattle, Washington, USA. AAAI. ISBN 1-57735-134-7. See <http://www.aaai.org/Conferences/IAAI/iaai01.php> [accessed 2007-09-06].
- [322] Jan Aikins and Howard Shrobe, editors. *Proceedings of the The Seventh Conference on Innovative Applications of Artificial Intelligence (IAAI 1995)*, August 21–31, 1995, Montréal, Québec, Canada. ISBN 978-0-929280-81-3. See <http://www.aaai.org/Conferences/IAAI/iaai95.php> [accessed 2007-09-06].
- [323] Elizabeth Byrnes and Jan Aikins, editors. *Proceedings of the The Sixth Conference on Innovative Applications of Artificial Intelligence (IAAI 1994)*, July 13–August 4, 1994, Convention Center, Seattle, Washington, USA. ISBN 978-0-929280-62-2. See <http://www.aaai.org/Conferences/IAAI/iaai94.php> [accessed 2007-09-06].
- [324] *Proceedings of the The Fifth Conference on Innovative Applications of Artificial Intelligence (IAAI 1993)*, July 11–15, 1993, Washington, DC, USA. AAAI. ISBN 0-929280-46-6. See <http://www.aaai.org/Conferences/IAAI/iaai93.php> [accessed 2007-09-06].
- [325] A. Carlisle Scott and Philip Klahr, editors. *Proceedings of the The Fourth Conference on Innovative Applications of Artificial Intelligence (IAAI-92)*, July 12–16, 1992, San Jose, California, USA. AAAI. ISBN 0-262-69155-8. See <http://www.aaai.org/Conferences/IAAI/iaai912.php> [accessed 2007-09-06].
- [326] Reid G. Smith and A. Carlisle Scott, editors. *Proceedings of the The Third Conference on Innovative Applications of Artificial Intelligence (IAAI-91)*, June 14–19, 1991, Anaheim, California, USA. AAAI. ISBN 0-262-69148-5. See <http://www.aaai.org/Conferences/IAAI/iaai91.php> [accessed 2007-09-06].
- [327] Alain T. Rappaport and Reid G. Smith, editors. *Proceedings of the The Second Conference on Innovative Applications of Artificial Intelligence (IAAI-90)*, May 1–3,

- 1990, Washington, DC, USA. AAAI. ISBN 0-262-68068-8. Published in 1991. See <http://www.aaai.org/Conferences/IAAI/iaai90.php> [accessed 2007-09-06].
- [328] Herb Schorr and Alain T. Rappaport, editors. *Proceedings of the The First Conference on Innovative Applications of Artificial Intelligence (IAAI-89)*, March 28–30, 1989, Stanford University, Stanford, California, USA. AAAI. ISBN 978-0-262-69137-6. Published in 1991. See <http://www.aaai.org/Conferences/IAAI/iaai89.php> [accessed 2007-09-06].
- [329] Bruno Apolloni, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of 11th International Conference on the Knowledge-Based Intelligent Information and Engineering Systems, KES 2007, and the XVII Italian Workshop on Neural Networks, Part I*, volume 4692 of *Lecture Notes in Computer Science (LNCS)*, September 12–14, 2007, Vietri sul Mare, Italy. Springer. ISBN 978-3-540-74817-5. See also [330, 331].
- [330] Bruno Apolloni, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of 11th International Conference on the Knowledge-Based Intelligent Information and Engineering Systems, KES 2007, and the XVII Italian Workshop on Neural Networks, Part II*, volume 4693 of *Lecture Notes in Computer Science (LNCS)*, September 12–14, 2007, Vietri sul Mare, Italy. Springer. ISBN 978-3-540-74826-7. See also [329, 331].
- [331] Bruno Apolloni, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of 11th International Conference on the Knowledge-Based Intelligent Information and Engineering Systems, KES 2007, and the XVII Italian Workshop on Neural Networks, Part III*, volume 4694 of *Lecture Notes in Computer Science (LNCS)*, September 12–14, 2007, Vietri sul Mare, Italy. Springer. ISBN 978-3-540-74828-1. See also [329, 330].
- [332] Bogdan Gabrys, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 10th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2006, Part I*, volume 4251 of *Lecture Notes in Computer Science (LNCS)*, October 9–11, 2006, Bournemouth, UK. Springer. ISBN 3-540-46535-9. See also [333, 334].
- [333] Bogdan Gabrys, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 10th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2006, Part II*, volume 4252 of *Lecture Notes in Computer Science (LNCS)*, October 9–11, 2006, Bournemouth, UK. Springer. ISBN 3-540-46537-5. See also [332, 334].
- [334] Bogdan Gabrys, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 10th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2006, Part III*, volume 4253 of *Lecture Notes in Computer Science (LNCS)*, October 9–11, 2006, Bournemouth, UK. Springer. ISBN 3-540-46542-1. See also [332, 333].
- [335] Rajiv Khosla, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 9th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2005, Part I*, volume 3681 of *Lecture Notes in Computer Science (LNCS)*, September 14–16, 2005, Melbourne, Australia. Springer. ISBN 3-540-28894-5. See also [336, 337, 338].
- [336] Rajiv Khosla, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 9th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2005, Part II*, volume 3682 of *Lecture Notes in Computer Science (LNCS)*, September 14–16, 2005, Melbourne, Australia. Springer. ISBN 3-540-28895-3. See also [335, 337, 338].
- [337] Rajiv Khosla, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 9th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2005, Part III*, volume 3683 of *Lecture Notes in Computer Science (LNCS)*, September 14–16, 2005, Melbourne, Australia. Springer. ISBN 3-540-28896-1. See also [335, 337, 338].

- [338] Rajiv Khosla, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 9th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2005, Part IV*, volume 3684 of *Lecture Notes in Computer Science (LNCS)*, September 14–16, 2005, Melbourne, Australia. Springer. ISBN 3-540-28897-X. See also [335, 336, 337].
- [339] Mircea Gh. Negoita, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 8th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2004, Part I*, volume 3213 of *Lecture Notes in Computer Science (LNCS)*, September 20–25, 2004, Wellington, New Zealand. Springer. ISBN 3-540-23318-0. See <http://www.kesinternational.org/kes2004/> [accessed 2008-07-27]. See also [340, 341].
- [340] Mircea Gh. Negoita, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 8th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2004, Part II*, volume 3214 of *Lecture Notes in Computer Science (LNCS)*, September 20–25, 2004, Wellington, New Zealand. Springer. ISBN 3-540-23206-0. See <http://www.kesinternational.org/kes2004/> [accessed 2008-07-27]. See also [339, 341].
- [341] Mircea Gh. Negoita, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 8th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2004, Part III*, volume 3215 of *Lecture Notes in Computer Science (LNCS)*, September 20–25, 2004, Wellington, New Zealand. Springer. ISBN 3-540-23205-2. See <http://www.kesinternational.org/kes2004/> [accessed 2008-07-27]. See also [339, 340].
- [342] Vasile Palade, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 7th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2003, Part I*, volume 2773 of *Lecture Notes in Computer Science (LNCS)*, September 3–5, 2003, Oxford, UK. Springer. ISBN 3-540-40803-7. See also [343].
- [343] Vasile Palade, Robert J. Howlett, and Lakhmi C. Jain, editors. *Proceedings of the 7th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2003, Part II*, volume 2774 of *Lecture Notes in Computer Science (LNCS)*, September 3–5, 2003, Oxford, UK. Springer. ISBN 3-540-40804-5. See also [342].
- [344] E. Damiani, Lakhmi C. Jain, and Robert J. Howlett, editors. *Proceedings of the Sixth International Conference on Knowledge-Based Intelligent Information & Engineering Systems, KES 2002*, September 16–18, 2002, Podere d'Ombriano, Crema, Italy. See <http://www.bton.ac.uk/kes/kes2002/kes2002.html/> [accessed 2008-07-28].
- [345] Lakhmi C. Jain, N. Baba, and Robert J. Howlett, editors. *Proceedings of the Fifth International Conference on Knowledge-Based Intelligent Information Engineering Systems & Allied Technologies, KES 2001*, September 6–8, 2001, Osaka-Kyoiku University, Osaka, Japan and Nara-ken New Public Hall, Nara, Japan. See <http://www.bton.ac.uk/kes/kes2001/> [accessed 2008-07-28].
- [346] Robert J. Howlett and Lakhmi C. Jain, editors. *Proceedings of the Fourth International Conference on Knowledge-Based Intelligent Information Engineering Systems & Allied Technologies, KES 2000, Volume 1*, August 30–September 1, 2000, Brighton, UK. IEEE. ISBN 0-7803-6400-7. See also [347].
- [347] Robert J. Howlett and Lakhmi C. Jain, editors. *Proceedings of the Fourth International Conference on Knowledge-Based Intelligent Information Engineering Systems & Allied Technologies, KES 2000, Volume 2*, August 30–September 1, 2000, Brighton, UK. IEEE. ISBN 0-7803-6400-7. See also [346].
- [348] Lakhmi C. Jain, editor. *Proceedings of the Third International Conference on Knowledge-Based Intelligent Information Engineering Systems, KES 1999*, August 31–September 1, 1999, Adelaide, South Australia. IEEE. ISBN 0-7803-5578-4.

- [349] Lakhmi C. Jain and R. K. Jain, editors. *Proceedings of the 2nd International Conference on Knowledge-Based Intelligent Electronic Systems, KES 1998, Part I*, April 21–23, 1998, Adelaide, South Australia. IEEE. See also [350, 351].
- [350] Lakhmi C. Jain and R. K. Jain, editors. *Proceedings of the 2nd International Conference on Knowledge-Based Intelligent Electronic Systems, KES 1998, Part II*, April 21–23, 1998, Adelaide, South Australia. IEEE. See also [349, 351].
- [351] Lakhmi C. Jain and R. K. Jain, editors. *Proceedings of the 2nd International Conference on Knowledge-Based Intelligent Electronic Systems, KES 1998, Part III*, April 21–23, 1998, Adelaide, South Australia. IEEE. See also [349, 350].
- [352] Lakhmi C. Jain, editor. *Proceedings of the 1st International Conference on Knowledge-Based Intelligent Electronic Systems, KES 1997, Part I*, May 21–23, 1997, Adelaide, South Australia. IEEE. ISBN 0-7803-3755-7. See also [353].
- [353] Lakhmi C. Jain, editor. *Proceedings of the 1st International Conference on Knowledge-Based Intelligent Electronic Systems, KES 1997, Part II*, May 21–23, 1997, Adelaide, South Australia. IEEE. ISBN 0-7803-3755-7. See also [352].
- [354] Matthias Ehrgott, editor. *Proceedings of the 19th International Conference on Multiple Criteria Decision Making: MCDM for Sustainable Energy and Transportation Systems (MCDM'2008)*, June 7–12, 2008, Auckland, New Zealand. See <https://secure.orsnz.org.nz/mcdm2008/> [accessed 2007-09-10].
- [355] Constantin Zopounidis, editor. *Proceedings of the 18th International Conference on Multiple Criteria Decision Making (MCDM'2006)*, June 9–13, 2006, MAICH (Mediterranean Agronomic Institute of Chania) Conference Centre, Chania, Crete, Greece. See <http://www.dpem.tuc.gr/fel/mcdm2006/> [accessed 2007-09-10].
- [356] Bill Wedley, editor. *Proceedings of the 17th International Conference on Multiple Criteria Decision Making (MCDM'2004)*, August 6–9, 2004, Whistler Conference Center, Whistler, British Columbia, Canada. See <http://www.bus.sfu.ca/events/mcdm/> [accessed 2007-09-10].
- [357] Mikulas Luptacik and Rudolf Vetschera, editors. *Proceedings of the First MCDM Winter Conference, 16th International Conference on Multiple Criteria Decision Making (MCDM'2002)*, February 12–22, 2002, Hotel Panhans, Semmering, Austria. See <http://orgwww.bwl.univie.ac.at/mcdm2002> [accessed 2007-09-10].
- [358] Murat Köksalan and Stanley Zionts, editors. *Proceedings of the 15th International Conference on Multiple Criteria Decision Making: Multiple Criteria Decision Making in the New Millennium (MCDM'2000)*, volume 507 of *Lecture Notes in Economics and Mathematical Systems*, July 10–14, 2000, Middle East Technical University, Ankara, Turkey. Springer. ISBN 978-3540423775. See <http://mcdm2000.ie.metu.edu.tr/> [accessed 2007-09-10]. Published November 9, 2001.
- [359] Theo Stewart, editor. *Proceedings of the 13th International Conference on Multiple Criteria Decision Making: Trends in Multicriteria Decision Making (MCDM'1997)*, volume 465 of *Lecture Notes in Economics and Mathematical Systems*, January 6–10, 1997, University of Cape Town, Cape Town, South Africa. Springer-Verlag TeLOS. ISBN 978-3540647416. See <http://www.uct.ac.za/depts/math/mcdm97> [accessed 2007-09-10]. Published January 15, 1999.
- [360] Günter Fandel, Thomás Gál, and Thomas Hanne, editors. *Proceedings of the 12th International Conference on Multiple Criteria Decision Making (MCDM'1995)*, volume 448 of *Lecture Notes in Economics and Mathematical Systems*, 1995, Hagen, Germany. Springer. ISBN 978-3540620976. Published on March 21, 1997.
- [361] João Climaco, editor. *Proceedings of the 11th International Conference on Multiple Criteria Decision Making: Multicriteria Analysis (MCDM'1994)*, August 1–6, 1994, Coimbra, Portugal. Springer. ISBN 978-3540620747. Published in May 1997.
- [362] G. H. Tzeng, Herbert F. Wang, U. P. Wen, and Po-Lung Yu, editors. *Proceedings of the 10th International Conference on Multiple Criteria Decision Making: Expand*

- and *Enrich the Domains of Thinking and Application (MCDM'1992)*, 1992, Taipei, Taiwan. Springer. ISBN 978-0387942971. Published in June 1994.
- [363] Richard Soland, Ambrose Coicochea, L. Duckstein, and Stanley Zionts, editors. *Proceedings of the 9th International Conference on Multiple Criteria Decision Making: Theory and Applications in Business, Industry, and Government (MCDM'1990)*, 1990, Fairfax, USA. Springer. ISBN 978-0387978055. Published in July 1992.
- [364] A. Geoff Lockett and Gerd Islei, editors. *Proceedings of the 8th International Conference on Multiple Criteria Decision Making: Improving Decision Making in Organizations (MCDM'1988)*, Lecture Notes in Economics and Mathematical Systems, 1988, Manchester, UK. Springer. ISBN 978-0387517957. Published in December 1989.
- [365] H. Nakayama and Y. Sawaragi, editors. *Proceedings of the 7th International Conference on Multiple Criteria Decision Making (MCDM'1986)*, 1986, Kyoto, Japan.
- [366] Yacov Y. Haimes, editor. *Proceedings of the 6th International Conference on Multiple Criteria Decision Making: Decision Making with Multiple Objectives (MCDM'1984)*, volume 242 of *Lecture Notes in Economics and Mathematical Systems*, June 4–6, 1984, Cleveland, Ohio, USA. Springer-Verlag Berlin and Heidelberg GmbH & Co. KG. ISBN 978-3540152231. Published December 31, 1985.
- [367] Pierre Hansen, editor. *Proceedings of the 5th International Conference on Multiple Criteria Decision Making: Essays and Surveys on Multiple Criteria Decision Making (MCDM'1982)*, volume 209 of *Lecture Notes in Economics and Mathematical Systems*, 1982, Mons, Belgium. Springer. ISBN 978-0387119915. Published in March 1983.
- [368] Joel N. Morse, editor. *Proceedings of the 4th International Conference on Multiple Criteria Decision Making: Organizations, Multiple Agents With Multiple Criteria (MCDM'1980)*, Lecture Notes in Economics and Mathematical Systems, 1980, Newark, Delaware, USA. Springer. ISBN 978-0387108216. Published in July 1981.
- [369] Günter Fandel and Thomás Gál, editors. *Proceedings of the 3rd International Conference on Multiple Criteria Decision Making: Theory and Application (MCDM'1979)*, volume 17 of *Lecture Notes in Economics and Mathematical Systems*, August 20–24, 1979, Hagen/Königswinter, (West) Germany. Springer-Verlag, Berlin, Germany. ISBN 978-0387099637. Published in May 1980.
- [370] Stanley Zionts, editor. *Proceedings of the 2nd International Conference on Multiple Criteria Decision Making (MCDM'1977)*, 1977, Buffalo, New York, USA.
- [371] Herve Thiriez and Stanley Zionts, editors. *Proceedings of the 1st International Conference on Multiple Criteria Decision Making (MCDM'1975)*, Lecture notes in economics and mathematical systems, May 21–23, 1975, Jouy-en-Josas, France. Springer. ISBN 978-0387077949. Published in 1976.
- [372] *Proceedings of the 13th International Conference on Soft Computing, MENDEL'07*, September 5–7, 2007, Prague, Czech Republic.
- [373] *Proceedings of the 12th International Conference on Soft Computing, MENDEL'06*, May 31–June 2, 2006, Brno University of Technology, Brno, Czech Republic. ISBN 80-214-3195-4.
- [374] *Proceedings of the 11th International Conference on Soft Computing, MENDEL'05*, June 15–17, 2005, Brno University of Technology, Brno, Brno, Czech Republic. ISBN 80-214-2961-5.
- [375] *Proceedings of the 10th International Conference on Soft Computing, MENDEL'04*, June 16–18, 2004, Brno University of Technology Faculty of Mechanical Engineering, Brno, Brno, Czech Republic. ISBN 80-214-2676-4.
- [376] *Proceedings of the 9th International Conference on Soft Computing, MENDEL'03*, 2003, Brno University of Technology, Brno, Czech Republic. ISBN 80-214-2411-7.
- [377] *Proceedings of the 8th International Conference on Soft Computing, MENDEL'02*, June 5–7, 2002, Brno University of Technology, Brno, Czech Republic. ISBN 80-214-2135-5.

- [378] *Proceedings of the 7th International Conference on Soft Computing, Evolutionary Computation, Genetic Programming, Fuzzy Logic, Rough Sets, Neural Networks, Fractals, Bayesian Methods, MENDEL 2001*, June 6–8, 2001, Brno University of Technology, Brno, Czech Republic. ISBN 80-7248-107-X.
- [379] Pavel Osmera, editor. *Proceedings of the 6th International Conference on Soft Computing, MENDEL 2000*, June 7–9, 2000, Brno University of Technology, Brno, Czech Republic. ISBN 80-214-1609-2.
- [380] *Proceedings of the 3rd International Conference on Genetic Algorithms*, June 1997, Brno University of Technology, Brno, Czech Republic. ISBN 80-214-1131-7.
- [381] *Proceedings of the 4th International Mendel Conference on Genetic Algorithms, Optimisation Problems, Fuzzy Logic, Neural Networks, Rough Sets*, June 24–26, 1998, Technical University of Brno, Faculty of Mechanical Engineering, Brno, Czech Republic. ISBN 80-214-1131-7.
- [382] *Proceedings of the 2nd International Mendel Conference on Genetic Algorithms, MENDEL 1997*, June 26–28, 1997, Brno University of Technology, Brno, Czech Republic. ISBN 80-214-0769-7.
- [383] *Proceedings of the 2nd International Mendel Conference on Genetic Algorithms, MENDEL 1996*, June 26–28, 1996, Brno University of Technology, Brno, Czech Republic. ISBN 80-214-0769-7.
- [384] *Proceedings of Mendel'95, the 1st International Mendel Conference on Genetic Algorithms*, September 1995, Brno, Czech Republic.
- [385] *The Seventh Metaheuristics International Conference*, June 25–29, 2007, Montréal, Québec, Canada. See <http://www.mic2007.ca/> [accessed 2007-09-12].
- [386] *6th Metaheuristics International Conference (MIC2005)*, August 22–26, 2005, Vienna, Austria. See <http://www.mic2005.org/> [accessed 2007-09-12].
- [387] Toshihide Ibaraki, Koji Nonobe, and Mutsunori Yagiura, editors. *Fifth Metaheuristics International Conference (MIC2003) – Metaheuristics: Progress as Real Problem Solvers*, volume 32 of *Operations Research/Computer Science Interfaces*, August 25–28, 2003, Kyoto International Conference Hall, Kyoto, Japan. Springer, Berlin Heidelberg New York. ISBN 0-387-25382-3. Published in 2005. See <http://www-or.amp.i.kyoto-u.ac.jp/mic2003/> [accessed 2007-09-12].
- [388] Mauricio G. C. Resende and Jorge Pinho de Sousa, editors. *4th Metaheuristics International Conference (MIC2001) – Metaheuristics: Computer Decision-Making*, volume 86 of *Applied Optimization*, July 16–20, 2001, Porto, Portugal. Kluwer Academic Publishers, Boston, USA. ISBN 1-4020-7653-3. Published in November 30, 2003. See <http://paginas.fe.up.pt/~gauti/MIC2001/> [accessed 2007-09-12].
- [389] Celso C. Ribeiro and Pierre Hansen, editors. *3rd Metaheuristics International Conference (MIC'99) – Essays and Surveys in Metaheuristics*, volume 15 of *Operations Research/Computer Science Interfaces*, July 19–23, 1999, Angra dos Reis, Brazil. Kluwer Academic Publishers, Boston, USA. ISBN 0-7923-7520-3. Published in 2001.
- [390] Stefan Voss, Silvano Martello, Ibrahim H. Osman, and Cathérine Roucairol, editors. *2nd Metaheuristics International Conference (MIC'97) – Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*, July 21–24, 1997, Sophia Antipolis, France. Kluwer Academic Publishers, Boston, USA. ISBN 0-7923-8369-9. Published in 1998.
- [391] Ibrahim H. Osman and James P. Kelly, editors. *1st Metaheuristics International Conference (MIC'95) – Meta-Heuristics: The Theory and Applications*, July 22–26, 1995, Breckenridge, Colorado, USA. Kluwer Academic Publishers, Boston, USA. ISBN 0-7923-9700-2. Published in 1996.
- [392] Alexander F. Gelbukh and Angel Fernando Kuri Morales, editors. *Proceedings of the MICA 2007: Advances in Artificial Intelligence, 6th Mexican International Conference on Artificial Intelligence*, volume 4827/2007 of *Lecture Notes in Computer Science (LNCS)*, Subseries *Lecture Notes in Artificial Intelligence (LNAI)*, Novem-

- ber 4–10, 2007, Aguascalientes, México. Springer Berlin / Heidelberg. ISBN 978-3-540-76630-8. doi:10.1007/978-3-540-76631-5.
- [393] Alexander F. Gelbukh and Carlos A. Reyes García, editors. *Proceedings of the MICAI 2006: Advances in Artificial Intelligence, 5th Mexican International Conference on Artificial Intelligence*, volume 4293/2006 of *Lecture Notes in Computer Science (LNCS)*, Subseries *Lecture Notes in Artificial Intelligence (LNAI)*, November 13–17, 2006, Apizaco, México. Springer Berlin / Heidelberg. ISBN 3-540-49026-4, 978-3-540-49026-5. doi:10.1007/11925231. See also [394].
- [394] Rogelio Dávila, Mauricio Osorio, and Claudia Zepeda, editors. *Proceedings of the LoLaCOM06 Workshop on Logic, Language and Computation*, volume 220 of *CEUR Workshop Proceedings*, November 13–14, 2006, Instituto Tecnológico de Apizaco, Apizaco, Tlaxcala, México. CEUR-WS.org. Online available at <http://CEUR-WS.org/Vol-220/> and <ftp://sunsite.informatik.rwth-aachen.de/pub/publications/CEUR-WS/Vol-220.zip> [accessed 2008-06-29]. See also [393].
- [395] Alexander F. Gelbukh, Alvaro de Albornoz, and Hugo Terashima-Marín, editors. *Proceedings of the MICAI 2005: Advances in Artificial Intelligence, 4th Mexican International Conference on Artificial Intelligence*, volume 3789/2005 of *Lecture Notes in Computer Science (LNCS)*, Subseries *Lecture Notes in Artificial Intelligence (LNAI)*, November 14–18, 2005, Monterrey, México. Springer Berlin / Heidelberg. ISBN 3-540-29896-7, 978-3-540-29896-0. doi:10.1007/11579427.
- [396] Raul Monroy, Gustavo Arroyo-Figueroa, Luis Enrique Sucar, and Juan Humberto Sossa Azuela, editors. *Proceedings of the MICAI 2004: Advances in Artificial Intelligence, Third Mexican International Conference on Artificial Intelligence*, volume 2972/2004 of *Lecture Notes in Computer Science (LNCS)*, Subseries *Lecture Notes in Artificial Intelligence (LNAI)*, April 26–30 2004, Mexico City, México. Springer Berlin / Heidelberg. ISBN 3-540-21459-3, 978-3-540-21459-5. doi:10.1007/b96521.
- [397] Carlos Artemio Coello Coello, Alvaro de Albornoz, Luis Enrique Sucar, and Osvaldo Cairó Battistutti, editors. *Proceedings of the MICAI 2002: Advances in Artificial Intelligence, Second Mexican International Conference on Artificial Intelligence*, volume 2313/2002 of *Lecture Notes in Computer Science (LNCS)*, Subseries *Lecture Notes in Artificial Intelligence (LNAI)*, April 22–26 2002, Mérida, Yucatán, México. Springer Berlin / Heidelberg. ISBN 3-540-43475-5, 978-3-540-43475-7. doi:10.1007/3-540-46016-0.
- [398] Osvaldo Cairó, Luis Enrique Sucar, and Francisco J. Cantu, editors. *Proceedings of the MICAI 2000: Advances in Artificial Intelligence, Mexican International Conference on Artificial Intelligence*, volume 1793/2000 of *Lecture Notes in Computer Science (LNCS)*, Subseries *Lecture Notes in Artificial Intelligence (LNAI)*, April 11–14 2000, Acapulco, México. Springer Berlin / Heidelberg. ISBN 3-540-67354-7, 978-3-540-67354-5. doi:10.1007/10720076.
- [399] Xavier Gandibleux, Marc Sevaux, Kenneth Sörensen, and Vincent T'kindt, editors. *Metaheuristics for Multiobjective Optimisation*, volume 535 of *Lecture Notes in Economics and Mathematical Systems*. Springer, Berlin, January 2004. ISBN 354020637X, 978-3540206378.
- [400] David Corne, Marco Dorigo, Fred Glover, Dipankar Dasgupta, Pablo Moscato, Riccardo Poli, and Kenneth V. Price, editors. *New Ideas in Optimization*. McGraw-Hill's Advanced Topics In Computer Science Series. McGraw-Hill Ltd., UK, Maidenhead, UK, England, May 1999. ISBN 0-07-709506-5, 978-0077095062.
- [401] Lakhmi C. Jain and Janusz Kacprzyk, editors. *New Learning Paradigms in Soft Computing*, volume 84 (PR-2001-06) of *Studies in fuzziness and soft computing*. Physica-Verlag (A Springer-Verlag Company), Heidelberg, Germany, February 2002. ISBN 9783790814361. Partly online available at <http://books.google.de/books?hl=de&id=u2Ncyk0hYYwC> [accessed 2008-05-31].

- [402] Ashutosh Tiwari, Joshua Knowles, Erel Avineri, Keshav Dahal, and Rajkumar Roy, editors. *Applications of Soft Computing – Recent Trends*, volume 36 of *Advances in Soft Computing*. Springer-Verlag Berlin and Heidelberg GmbH & Co. KG, June 7, 2006. ISBN 978-3-540-29123-7. Editor-in-chief: Janusz Kacprzyk.
- [403] Patrick Siarry and Zbigniew Michalewicz, editors. *Advances in Metaheuristics for Hard Optimization*. Natural Computing Series. Springer-Verlag, LLC, New York, January 2008. ISBN 9783540729594. Partly online available at <http://books.google.de/books?id=TkMR1b-VCR4C> [accessed 2008-08-01]. Series editors: G. Rozenberg, Thomas Bäck, Agoston E. Eiben, J.N. Kok, and H.P. Spaink.
- [404] Godfrey C. Onwubolu and B. V. Babu, editors. *New Optimization Techniques in Engineering*, volume 141 of *Studies in Fuzziness and Soft Computing*. Springer-Verlag, Berlin/Heidelberg, Germany, March 5, 2004. ISBN 354020167X, 978-3540201670. Partly online available at <http://books.google.de/books?id=YRsQ-Pz0kAAC> [accessed 2008-06-07].
- [405] Panos M. Pardalos and Ding-Zhu Du. *Handbook of Combinatorial Optimization*, volume 1–3 of *Combinatorial Optimization*. Kluwer Academic Publishers, Springer Netherlands, October 31, 1998. ISBN 0792350197, 978-0792350194.
- [406] Colin R. Reeves, editor. *Modern Heuristic Techniques for Combinatorial Problems*. John Wiley & Sons, Inc., Blackwell Scientific Publishing, Halsted Press, New York, NY, USA / Oxford, England, April 1993. ISBN 0-470-22079-1, 978-0470220795.
- [407] Erricos John Kontoghiorghe, editor. *Handbook of Parallel Computing and Statistics*. Statistics: Textbooks and Monographs. Chapman & Hall / CRC Press / Marcel Dekker Inc, December 21, 2005. ISBN 082474067X, 9780824740672. Partly online available at <http://books.google.de/books?id=BnNnKPkFH2kC> [accessed 2008-08-03].
- [408] Illogic. One Brick. In Aesop Rock, editor, *Labor Days*, Aesop Rock, Omega One, and Blockhead, producers. Definitive Jux, September 18, 2001.
- [409] Thomas Bäck. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, January 1996. ISBN 0195099710.
- [410] Thomas Bäck, David B. Fogel, and Zbigniew Michalewicz, editors. *Handbook of Evolutionary Computation*. Computational Intelligence Library. Oxford University Press in cooperation with the Institute of Physics Publishing, Bristol, New York, ringbound edition, April 1997. ISBN 0750303921.
- [411] Thomas Bäck, Ulrich Hammel, and Hans-Paul Schwefel. Evolutionary computation: comments on the history and current state. *IEEE Transactions on Evolutionary Computation*, 1(1):3–17, April 1997. Online available at <http://sci2s.ugr.es/docencia/doctobio/EC-History-IEEE-TEC-1-1-1997.pdf> and <http://citeseer.ist.psu.edu/601414.html> [accessed 2007-08-24].
- [412] Ernst Mayr. *The Growth of Biological Thought*. Harvard University Press, 1982, Cambridge, MA, USA. ISBN 0-674-36446-5.
- [413] J. Koen van der Hauw. Evaluating and improving steady state evolutionary algorithms on constraint satisfaction problems. Master’s thesis, Computer Science Department of Leiden University, August 1996, Computer Science Department of Leiden University. supervisors: Gusztai Eiben and Han La Poutré. Online available at <http://citeseer.ist.psu.edu/128231.html> [accessed 2007-08-24].
- [414] Carlos Artemio Ceollos Coello. A Comprehensive Survey of Evolutionary-Based Multiobjective Optimization Techniques. *Knowledge and Information Systems*, 1(3):269–308, August 1999. Online available at <http://www.lania.mx/~ccoello/EMOO/informationfinal.ps.gz> and <http://citeseer.ist.psu.edu/coello98comprehensive.html> [accessed 2007-08-25].
- [415] David A. van Veldhuizen. *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*. PhD thesis, Air Force Institute of Technology, Air University, Wright-Patterson Air Force Base, Ohio, May 1999. AFIT/DS/ENG/99-

01. Online available at <http://handle.dtic.mil/100.2/ADA364478> and <http://citeseer.ist.psu.edu/vanveldhuizen99multiobjective.html> [accessed 2007-08-17].
- [416] Mark Ridley. *Evolution*. Blackwell Publishing Limited, eighth (october 1, 2003) edition, 1996. ISBN 978-1405103459.
- [417] John J. Greffentette and James E. Baker. How Genetic Algorithms Work: A Critical Look at Implicit Parallelism. In *Proceedings of the third International Conference on Genetic Algorithms ICGA*, pages 20–27, 1989. In proceedings [652].
- [418] John J. Grefenstette. Conditions for Implicit Parallelism. In *Foundations of genetic algorithms*, pages 252–261, 1990. In proceedings [188]. Online available at <http://citeseer.ist.psu.edu/41022.html> [accessed 2007-07-29].
- [419] Alberto Bertoni and Marco Dorigo. Implicit Parallelism in Genetic Algorithms. *Artificial Intelligence*, 61(2):307–314, 1993. Online available at <http://citeseer.ist.psu.edu/bertoni93implicit.html> [accessed 2007-07-29].
- [420] Michael D. Vose and Alden H. Wright. Form Invariance and Implicit Parallelism. *Evolutionary Computation*, 9(3):355–370, 2001. ISSN 1063-6560. doi:10.1162/106365601750406037. Online available at <http://citeseer.ist.psu.edu/610097.html> [accessed 2007-07-29].
- [421] Oliver John Sharpe. *Towards a Rational Methodology for Using Evolutionary Search Algorithms*. PhD thesis, University of Sussex, January 2000. Online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.36.1667> [accessed 2008-08-10].
- [422] Kenneth Alan De Jong, David B. Fogel, and Hans-Paul Schwefel. A History of Evolutionary Computation. In *Evolutionary Computation 1: Basic Algorithms and Operators*, chapter 6, pages 40–. Institute of Physics Publishing (IOP), 2000. Online available at <http://books.google.de/books?id=4HMYCq9US78C> [accessed 2008-06-11].
- [423] George E. P. Box from the Imperial Chemical Industry Limited, Dyestuffs Division Headquarters, Manchester. Evolutionary Operation: A Method for Increasing Industrial Productivity. *Applied Statistics*, 6(2):81–101, June 1957. ISSN 00359254.
- [424] George E. P. Box and Norman R. Draper. *Evolutionary operation. A statistical method for process improvement*. Wiley Publication in Applied Statistics. John Wiley & Sons, 1969. ISBN 047109305X, 978-0471093053.
- [425] W. Spendley, G. R. Hext, and F. R. Himsforth from the Imperial Chemical Industries Ltd., Billingham, Co. Durham, England. Sequential application of simplex designs in optimisation and evolutionary operation. *Technometrics*, 4(4):441–461, November 1962.
- [426] John Ashworth Nelder and Roger A. Mead. A Simplex Method for function Minimization. *Computer Journal*, 7:308–313, 1965.
- [427] Robert Michael Lewis, Virginia Joanne Torczon, and Michael W. Trosset. Direct Search Methods: Then and Now. Technical Report NASA/CR-2000-210125 and ICASE Report No. 2000-26, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton, VA, USA (Operated by Universities Space Research Association), National Aeronautics and Space Administration Langley Research Center Hampton, Virginia 23681-2199, May 2000. Online available at <http://historical.ncstrl.org/tr/pdf/icase/TR-2000-26.pdf> [accessed 2008-06-14].
- [428] F. E. Satterthwaite from the Statistical Engineering Institute, Wellesley, Massachusetts, USA and Merriam College, North Andover, Massachusetts, USA. Random Balance Experimentation. *Technometrics*, 1(2):111–137, May 1959.
- [429] F. E. Satterthwaite. REVOP or Random Evolutionary Operation. Merrimack College Technical Report 10-10-59, Merriam College, North Andover, Massachusetts, USA, 1959.
- [430] Carlos Artemio Coello Coello. A Short Tutorial on Evolutionary Multiobjective Optimization. In *First International Conference on Evolutionary Multi-Criterion*

- Optimization*, pages 21–40, 2001. In proceedings [521]. Online available at <http://citeseer.ist.psu.edu/coellocoello01short.html> [accessed 2007-07-29].
- [431] Nadia Nedjah, Enrique Alba, and Luiza De Macedo Mourelle, editors. *Parallel Evolutionary Computations*, volume 22/2006 of *Studies in Computational Intelligence*. Springer Berlin / Heidelberg, June 2006. ISBN 9783540328377. doi:10.1007/3-540-32839-4.
- [432] Masaya Shinkai, Hern'an Aguirre, and Kiyoshi Tanaka. Mutation strategy improves GAs performance on epistatic problems. In *Proceedings of the 2002 Congress on Evolutionary Computation, CEC'02*, volume 1, pages 968–973, 2002. doi:10.1109/CEC.2002.1007056. In proceedings [499].
- [433] Adam Prügel-Bennett and Alex Rogers. Modelling GA Dynamics. In *Theoretical Aspects of Evolutionary Computing*, pages 59–86. Springer, 2001. Original: 1999, Online available at <http://eprints.ecs.soton.ac.uk/13205/> [accessed 2007-07-28]. In collection [615].
- [434] Hsien-Chung Wu from the National Chi Nan University, Puli, Nantou 545, Taiwan. *Evolutionary Computation*. Department of Mathematics, National Kaohsiung Normal University, Kaohsiung 802, Taiwan, February 2005. Lecture notes. Online available at <http://nknucc.nknu.edu.tw/~hcwu/pdf/evolec.pdf> [accessed 2007-07-16].
- [435] Thomas Bäck and Hans-Paul Schwefel. An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation*, 1(1):1–23, Spring 1993. ISSN 1063-6560.
- [436] Gareth Jones. Genetic and Evolutionary Algorithms. In Paul von Ragué Schleyer and Johnny Gasteiger, editors, *Encyclopedia of Computational Chemistry*, volume III - Databases and Expert Systems, chapter 29. John Wiley & Sons, Ltd., September 1998. ISBN 978-0471965886. Online available at <http://www.wiley.com/legacy/wileychi/ecc/samples/sample10.pdf> [accessed 2007-08-17].
- [437] Claudio Rossi, Elena Marchiori, and Joost N. Kok. An adaptive evolutionary algorithm for the satisfiability problem. In *SAC'00: Proceedings of the 2000 ACM symposium on Applied computing*, volume 1, pages 463–469, 2000, Como, Italy. ACM Press, New York, NY, USA. ISBN 1-58113-240-9. doi:10.1145/335603.335912. Online available at <http://doi.acm.org/10.1145/335603.335912> and <http://citeseer.ist.psu.edu/335101.html> [accessed 2007-08-24].
- [438] Frank Hoffmeister and Thomas Bäck. Genetic Algorithms and Evolution Strategies – Similarities and Differences. In *PPSN I: Proceedings of the 1st Workshop on Parallel Problem Solving from Nature*, pages 455–469, 1990. Published 1991. In proceedings [592] and also [439].
- [439] Frank Hoffmeister and Thomas Bäck. Genetic Algorithms and Evolution Strategies – Similarities and Differences. Technical Report SYS-1/92, University of Dortmund - Systems Analysis, 1992. See also [438].
- [440] Hans-Paul Schwefel. *Evolution and Optimum Seeking: The Sixth Generation*. John Wiley & Sons, Inc., Wiley Interscience, New York, NY, USA, 1993. ISBN 0471571482, Har/Dsk: 978-0471571483. Har/Dsk edition, January 1995.
- [441] Alex Rogers and Adam Prügel-Bennett from the ISIS Research Group, Department of Electronics and Computer Science, University of Southampton, England. Modelling the Dynamics of a steady-state Genetic Algorithm. In *Foundations of Genetic Algorithms 5*, pages 57–68, 1998. In proceedings [735]. Online available at <http://eprints.ecs.soton.ac.uk/451/02/FOGA.ps> and <http://citeseer.ist.psu.edu/rogers99modelling.html> [accessed 2007-08-28].
- [442] Brian D. Davison and Khaled Rasheed from the Rutgers, The State University of New Jersey, New Brunswick, NJ. Effect of global parallelism on a steady state GA. In Erick Cantu-Paz and Bill Punch, editors, *Evolutionary Computation and Parallel Processing*, pages 167–170, 1999. Orlando, Florida, USA. Workshop part of GECCO 1999, see proceedings [572]. Online available at <http://citeseer.ist.psu.edu/>

- davison99effect.html [accessed <http://www.cse.lehigh.edu/~brian/pubs/1999/cec99/pgado.pdf>]2007-08-28.
- [443] David Noever and Subbiah Baskaran. Steady-State vs. Generational Genetic Algorithms: A Comparison of Time Complexity and Convergence Properties. Technical Report 92-07-032, Santa Fe Institute, 1992.
- [444] Deepti Chafekar, Jiang Xuan, and Khaled Rasheed. Constrained Multi-objective Optimization Using Steady State Genetic Algorithms. In *Proceedings of Genetic and Evolutionary Computation GECCO 2003*, pages 813–824, 2003. In proceedings [561]. Online available at <http://citeseer.ist.psu.edu/591279.html> and <http://www.cs.uga.edu/~khaled/papers/385.pdf> [accessed 2007-07-28].
- [445] Gilbert Syswerda. A Study of Reproduction in Generational and Steady State Genetic Algorithms. In *Proceedings of the First Workshop on Foundations of Genetic Algorithms FOGA*, pages 94–101, July 1990. In proceedings [188].
- [446] L. Darrell Whitley. The GENITOR Algorithm and Selective Pressure: Why Rank-Based Allocation of Reproductive Trials is Best. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 116–121, 1989. In proceedings [652]. Online available at <http://citeseer.ist.psu.edu/531140.html> and <http://www.cs.colostate.edu/~genitor/1989/ranking89.ps.gz> [accessed 2007-08-21].
- [447] Josh Jones and Terry Soule. Comparing genetic robustness in generational vs. steady state evolutionary algorithms. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 143–150, 2006. In proceedings [553]. Online available at <http://doi.acm.org/10.1145/1143997.1144024> [accessed 2007-08-24].
- [448] Bastien Chevreux. Genetische Algorithmen zur Molekülstrukturoptimierung. Master's thesis, Universität Heidelberg/Fachhochschule Heilbronn, Deutsches Krebsforschungszentrum Heidelberg, July 2001. Online available at <http://chevreux.org/diplom/diplom.html> [accessed 2007-07-29].
- [449] David E. Goldberg and Kalyanmoy Deb. A Comparative Analysis of Selection Schemes Used in Genetic Algorithms. In *Proceedings of the First Workshop on Foundations of Genetic Algorithms FOGA*, pages 69–93, 1990. In proceedings [188]. Online available at <http://www.cse.unr.edu/~sushil/class/gas/papers/Select.pdf> [accessed 2007-08-24].
- [450] Marco Laumanns, Eckart Zitzler, and Lothar Thiele. A Unified Model for Multi-Objective Evolutionary Algorithms with Elitism. In *Proceedings of the 2000 Congress on Evolutionary Computation CEC00*, pages 46–53, 2000. In proceedings [501]. Online available at <http://citeseer.ist.psu.edu/laumanns00unified.html> [accessed 2007-08-27].
- [451] Mihai Oltean. Evolving Evolutionary Algorithms for Function Optimization. In *Proceedings of the 5th International Workshop on Frontiers in Evolutionary Algorithms*, pages 295–298, 2003. In proceedings [545]. Online available at http://www.cs.ubbcluj.ro/~moltean/oltean_fea2003_1.pdf and <http://citeseer.ist.psu.edu/640208.html> [accessed 2007-08-24].
- [452] Mitchell A. Potter and Kenneth Alan De Jong. A Cooperative Coevolutionary Approach to Function Optimization. In *Parallel Problem Solving from Nature - PPSN III, International Conference on Evolutionary Computation. The Third Conference on Parallel Problem Solving from Nature*, pages 249–257, 1994. In proceedings [590]. Online available at <http://citeseer.ist.psu.edu/potter94cooperative.html> [accessed 2007-08-24].
- [453] Fernando Jiménez, José L. Verdegay, and Antonio F. Gómez-Skarmeta. Evolutionary techniques for constrained multiobjective optimization problems. In Kalyanmoy Deb, editor, *Multi-criterion Optimization Using Evolutionary Methods*, pages 115–116, 1999. In proceedings [572]. Online available at <http://citeseer.ist.psu.edu/208779.html> [accessed 2007-08-24].

- [454] Katharina Anna Lehmann and Michael Kaufmann. Evolutionary algorithms for the self-organized evolution of networks. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 563–570, 2005. doi:10.1145/1068009.1068105. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068105> and <http://www-pr.informatik.uni-tuebingen.de/mitarbeiter/katharinazweig/downloads/299-lehmann.pdf> [accessed 2008-05-28]. Lehmann is now known as Zweig.
- [455] Alex S. Fukunaga and Andre D. Stechert from the Jet Propulsion Laboratory, MS 525-3660, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109-8099. An evolutionary optimization system for spacecraft design. In *IEA/AIE'1997: Proceedings of the 10th international conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, pages 1–6. Goose Pond Press, 1997, Atlanta, Georgia. ISBN 905-6996-150. Online available at <http://alexf04.maclisp.org/Fukunaga-Stechert-IEA-1997.pdf> and <http://de.scientificcommons.org/17874369> [accessed 2007-08-24].
- [456] Kaisa Miettinen, Marko M. Mäkelä, Pekka Neittaanmäki, and Jacques Periaux, editors. *Evolutionary Algorithms in Engineering and Computer Science*. John Wiley & Sons, Ltd, Chichester, UK, July 16, 1999. ISBN 0471999024, 978-0471999027.
- [457] Shigeru Obayashi from the Department of Aeronautics and Space Engineering, Tohoku University, Sendai, 980-8579 Japan. Multidisciplinary Design Optimization of Aircraft Wing Planform Based on Evolutionary Algorithms. In *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE Press, October 11–14, 1998, La Jolla, California, USA. Online available at [accessed 2008-06-14].
- [458] Shu-Heng Chen, editor. *Evolutionary Computation in Economics and Finance*, volume 100 of *Studies in Fuzziness and Soft Computing*. Physica-Verlag Heidelberg, August 5, 2002. ISBN 3790814768.
- [459] Charles D. Day from the Internal Revenue Service, Statistics of Income Division. Application of an Evolutionary Algorithm to Multivariate Optimal Allocation in Stratified Sample Designs. In *SOI Tax Stats – Papers – 2006 American Statistical Association Conference, Joint Statistical Meeting 2006*. Tax Stats. Internal Revenue Service. United States Department of the Treasury, August 2006, Seattle, Washington. An Application of Genetic Algorithms to Multivariate Optimal Allocation in Stratified Sample Designs. Online available at <http://www.irs.gov/pub/irs-soi/06asaday.pdf> [accessed 2007-08-27].
- [460] Marco Aurelio Falcone, Heitor Silverio Lopes, and Leandro Dos Santos Coelho. Supply chain optimisation using evolutionary algorithms. *International Journal of Computer Applications in Technology*, 31(3/4):158–167, 2008. doi:10.1504/IJCAT.2008.018154.
- [461] Christopher D. Clack, editor. *Advanced Research Challenges in Financial Evolutionary Computing (ARC-FEC) Workshop 2008*, July 12, 2008, Renaissance Atlanta Hotel Downtown, 590 West Peachtree Street NW, Atlanta, Georgia 30308 USA. ACM Press, New York, NY, USA. ISBN 978-1-60558-131-6. Part of GECCO 2008, see also [550] and http://www.cs.ucl.ac.uk/financialcomputing/gecco_arcfec.htm [accessed 2008-07-21].
- [462] Gary B. Fogel and David W. Corne, editors. *Evolutionary Computation in Bioinformatics*. Academic Press. Elsevier LTD, Oxford, September 1, 2002. ISBN 1558607978, 978-1558607972. Partly online available at http://books.google.de/books?id=Jtj7Np09F_gC [accessed 2008-04-03].
- [463] Thomas Weise, Stefan Achler, Martin Göb, Christian Voigtmann, and Michael Zapf from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Evolving Classifiers – Evolutionary Algorithms in Data Mining. *Kasseler Informatikschriften (KIS)* 2007, 4, University of Kassel, University of Kassel, University of Kassel, September 28,

- 2007, University of Kassel. Persistent Identifier: urn:nbn:de:hebis:34-2007092819260. Online available at <http://www.it-weise.de/documents/files/WAGVZ2007DMC.pdf> and <http://kobra.bibliothek.uni-kassel.de/handle/urn:nbn:de:hebis:34-2007092819260> [accessed 2007-10-04].
- [464] Oscar Cordón, Francisco Herrera, and Luciano Sánchez. Evolutionary Learning Processes for Data Analysis in Electrical Engineering Applications. In D. Quagliarella, J. Périaux, C. Poloni, and G. Winter, editors, *Genetic Algorithms and Evolution Strategy in Engineering and Computer Science*, pages 205–224. John Wiley and Sons, Chichester, 1998. In collection [526]. Online available at <http://citeseer.ist.psu.edu/64737.html> and ftp://decsai.ugr.es/pub/arai/tech_rep/ga-fl/eurogen97.ps.Z [accessed 2007-08-25].
- [465] Ashish Ghosh and Lakhmi C. Jain, editors. *Evolutionary Computation in Data Mining*, volume 163/2005 of *Studies in Fuzziness and Soft Computing*. Springer Berlin / Heidelberg, 2005. ISBN 978-3-540-22370-2, 3-540-22370-3. doi:10.1007/3-540-32358-9.
- [466] Charles L. Karr and Eric Wilson. A Self-Tuning Evolutionary Algorithm Applied to an Inverse Partial Differential Equation. *Applied Intelligence*, 19(3):147–155, November 2003. ISSN 0924-669X (Print) 1573-7497 (Online). doi:10.1023/A:1026097605403. Online available at <http://www.springerlink.com/content/x7t26442h41q0221/fulltext.pdf> and <http://dx.doi.org/10.1023/A:1026097605403> [accessed 2007-08-27].
- [467] Swagatam Das, Amit Konar, and Uday K. Chakraborty. An efficient evolutionary algorithm applied to the design of two-dimensional IIR filters. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 2157–2163, 2005. doi:10.1145/1068009.1068364. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068364> [accessed 2007-08-27].
- [468] R. L. Johnston, editor. *Applications of Evolutionary Computation in Chemistry*, volume 110 of *Structure and Bonding*. Springer, Berlin, Germany, 2004. ISBN 978-3-540-40258-9, 3540402586. Partly online available at <http://books.google.de/books?id=AzbZIA5aT2oC> [accessed 2008-04-03].
- [469] Martin Damsbo, Brian S. Kinnear, Matthew R. Hartings, Peder T. Ruhoff, Martin F. Jarrold, and Mark A. Ratner. Application of evolutionary algorithm methods to polypeptide folding: Comparison with experimental results for unsolvated Ac-(Ala-Gly-Gly)₅-LysH⁺. *Proceedings of the National Academy of Science of the United States of America*, 101(19):7215–7222, May 2004. ISSN 1091-6490. Online available at <http://www.pnas.org/cgi/reprint/101/19/7215.pdf?ck=nck> and <http://www.pnas.org/content/vol101/issue19/> [accessed 2007-08-27].
- [470] Elena Marchiori and Adri G. Steenbeek. An Evolutionary Algorithm for Large Scale Set Covering Problems with Application to Airline Crew Scheduling. In *Proceedings of Real-World Applications of Evolutionary Computing, EvoWorkshops 2000*, pages 367–381, 2000. In proceedings [541]. Online available at <http://citeseer.ist.psu.edu/marchiori00evolutionary.html> [accessed 2007-08-27].
- [471] C. S. Chang and Chung Min Kwan from the Department of Electrical and Computer Engineering, National University of Singapore. Evaluation of evolutionary algorithms for multi-objective train schedule optimization. In *AI 2004: Advances in Artificial Intelligence*, volume 3339/2004 of *Lecture Notes in Artificial Intelligence, subseries of Lecture Notes in Computer Science (LNCS)*, pages 803–815. Springer-Verlag, 2004. ISBN 978-3-540-24059-4. doi:10.1007/b104336.
- [472] Chung Min Kwan and C. S. Chang from the Department of Electrical and Computer Engineering, National University of Singapore. Application of Evolutionary Algorithm on a Transportation Scheduling Problem – The Mass Rapid Transit. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation (CEC'2005)*, volume 2, pages 987–994, 2005. In proceedings [496]. Online available at <http://www.lania.mx/~ccoello/EM00/kwan05.pdf.gz> [accessed 2007-08-27].

- [473] Pascal Côté, Tony Wong, and Robert Sabourin from the Department of Automated Manufacturing Engineering, École de technologie supérieure, Université du Québec. Application of a Hybrid Multi-Objective Evolutionary Algorithm to the Uncapacitated Exam Proximity Problem. In Edmund K. Burke and Michael Trick, editors, *Proceedings of the 5th International Conference on Practice and Theory of Automated Timetabling (PATAT 2004)*, pages 151–168, 2004. Online available at <http://citeseer.ist.psu.edu/653221.html> and <http://www.asap.cs.nott.ac.uk/patat/patat04/151.pdf> [accessed 2007-08-27].
- [474] István Borgulya from the University of Pécs, Hungary. A Cluster-based Evolutionary Algorithm for the Single Machine Total Weighted Tardiness-scheduling Problem. *Journal of Computing and Information Technology - CIT*, 10(3):211–217, September 2002. Online available at <http://cit.zesoi.fer.hr/downloadPaper.php?paper=384> [accessed 2008-04-07].
- [475] Keigo Watanabe and M. M. A. Hashem. *Evolutionary Computations – New Algorithms and their Applications to Evolutionary Robots*. Studies in Fuzziness and Soft Computing. Springer, Berlin, Germany, March 2004. ISBN 3540209018, 978-3540209010. Partly online available at <http://books.google.de/books?id=G5HPNrJNTKsC> [accessed 2008-04-03].
- [476] Stefano Cagnoni, Evelynne Lutton, and Gustavo Olague, editors. *Genetic and Evolutionary Computation for Image Processing and Analysis*, volume 8 of *EURASIP Book Series on Signal Processing and Communications*. Hindawi Publishing Corporation, 410 Park Avenue, 15th Floor, #287 pmb, New York, NY 10022, USA, February 2008. ISBN 978-977-454-001-1. Online available at <http://www.hindawi.com/books/9789774540011.pdf> [accessed 2008-05-17].
- [477] Bertram Nickolay, Bernd Schneider, and Stefan Jacob from the Fraunhofer-IPK Berlin, Pascalstrae 8-9, 10587 Berlin, Germany. Parameter optimisation of an image processing system using evolutionary algorithms. In *Computer Analysis of Images and Patterns, Proceedings of the 7th International Conference, CAIP'97*, volume 1296/1997 of *Lecture Notes in Computer Science (LNCS)*, pages 637–644. Springer Berlin / Heidelberg, September 10-12 1997, Kiel, Germany. ISBN 978-3-540-63460-7. doi:10.1007/3-540-63460-6_173.
- [478] Mark C. Sinclair, David Corne, and George D. Smith, editors. *Proceedings of Evolutionary Telecommunications: Past, Present and Future – A Bird-of-a-feather Workshop at GECCO 99*, July 13, 1999, Orlando, Florida, USA. See also [572, 573]. Online available at <http://uk.geocities.com/markcsinclair/etppf.html> [accessed 2008-06-25].
- [479] Mark C. Sinclair, David Corne, and George D. Smith. Evolutionary Telecommunications: Past, Present and Future. In *Proceedings of Evolutionary Telecommunications: Past, Present and Future – A Bird-of-a-feather Workshop at GECCO 99*, 1999. In proceedings [478]. Online available at http://uk.geocities.com/markcsinclair/ps/etppf_sin_ws.ps.gz and <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.53.6326&rep=rep1&type=pdf> [accessed 2008-08-01].
- [480] Mark C. Sinclair from the Electronic Systems Engineering, University of Essex, Colchester CO4 3SQ, UK. Evolutionary Telecommunications: A Summary. In *Proceedings of Evolutionary Telecommunications: Past, Present and Future – A Bird-of-a-feather Workshop at GECCO 99*, 1999. In proceedings [478]. Online available at http://uk.geocities.com/markcsinclair/ps/etppf_sin_sum.ps.gz and <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.38.6797> [accessed 2008-08-01], presentation online available at http://uk.geocities.com/markcsinclair/ps/etppf_sin_slides.ps.gz [accessed 2008-08-01].
- [481] Pablo Cortés Achedad, Luis Onieva Giménez, Jesús Muñozuri Sanz, and José Guadix Martín from the Ingeniería de Organización. Escuela Técnica Superior de Ingenieros, University of Seville, 41092, Seville, Spain. A Revision of Evolutionary Computa-

- tion Techniques in Telecommunications and An Application for The Network Global Planning Problem. In *Success in Evolutionary Computation*, pages 239–262. Springer Berlin / Heidelberg, 2008. doi:10.1007/978-3-540-76286-7_11.
- [482] Tadashi Nakano and Tatsuya Suda from the Department of Information and Computer Science, University of California, Irvine, Irvine, CA 92697-3425, USA. Adaptive and Evolvable Network Services. In *Genetic and Evolutionary Computation – GECCO 2004*, pages 151–162, 2004. In proceedings [558].
- [483] Tadashi Nakano and Tatsuya Suda. Self-organizing network services with evolutionary adaptation. *IEEE Transactions on Neural Networks*, 16(5):1269–1278, September 2005. ISSN 1045-9227. doi:10.1109/TNN.2005.853421. Online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.74.9149> [accessed 2008-06-23].
- [484] Alejandro Quintero and Samuel Pierre from the Department of Computer Engineering, Mobile Computing, and Networking Research (LARIM), École Polytechnique de Montréal, C.P. 6079, Succ. Centre-Ville, Montréal, Québec, Canada H3C 3A7. Evolutionary approach to optimize the assignment of cells to switches in personal communication networks. *Computer Communications*, 26(9):927–938, June 2, 2003. doi:10.1016/S0140-3664(02)00238-4. Online available at [http://dx.doi.org/10.1016/S0140-3664\(02\)00238-4](http://dx.doi.org/10.1016/S0140-3664(02)00238-4) [http://dx.doi.org/10.1016/S0140-3664\(02\)00238-4](http://dx.doi.org/10.1016/S0140-3664(02)00238-4) [accessed 2008-08-01].
- [485] Enrique Alba, C. Cotta, J. Francisco Chicano, and A. J. Nebro. Parallel Evolutionary Algorithms in Telecommunications: Two Case Studies. In *Proceedings of the CACIC’02*, 2002, Buenos Aires, Argentina. Online available at <http://www.lcc.uma.es/~eat/.pdf/cacic02.pdf> and <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.8.4073> [accessed 2008-08-01].
- [486] Enrique Alba and J. Francisco Chicano from the Department Lenguajes y Ciencias de la Computación, ETSI Informática, University of Málaga, Campus de Teatinos, 29071 - Málaga - Spain. Evolutionary algorithms in telecommunications. In *IEEE Mediterranean Electrotechnical Conference, MELECON 2006*, pages 795–798, May 16-19, 2006. doi:10.1109/MELCON.2006.1653218. Online available at <http://dea.brunel.ac.uk/ugprojects/docs/melecon06.pdf> [accessed 2008-08-01].
- [487] David E. Clark, editor. *Evolutionary Algorithms in Molecular Design*, volume 8 of *Methods and Principles in Medicinal Chemistry*. Wiley-VCH, Weinheim and New York, September 11 2000. ISBN 3527301550, 978-3527301553. Series editors: Raimund Mannhold, Hugo Kubinyi, and Hendrik Timmerman.
- [488] Stephen Smith and Stefano Cagnoni, editors. *Medical Applications of Genetic and Evolutionary Computation (MedGEC) Workshop 2008*, July 12, 2008, Renaissance Atlanta Hotel Downtown, 590 West Peachtree Street NW, Atlanta, Georgia 30308 USA. ACM Press, New York, NY, USA. ISBN 978-1-60558-131-6. Part of GECCO 2008, see also [550] and <http://www.elec.york.ac.uk/intsys/events/MedGEC2008/home.htm> [accessed 2008-07-21].
- [489] Laurence D. Merkle and Frank Moore, editors. *Defense Applications of Computational Intelligence (DACI) Workshop 2008*, July 12, 2008, Renaissance Atlanta Hotel Downtown, 590 West Peachtree Street NW, Atlanta, Georgia 30308 USA. ACM Press, New York, NY, USA. ISBN 978-1-60558-131-6. Part of GECCO 2008, see also [550].
- [490] William Rand, Sevan G. Ficici, and Rick Riolo, editors. *Evolutionary Computation and Multi-Agent Systems and Simulation (ECoMASS) Workshop 2008*, July 12, 2008, Renaissance Atlanta Hotel Downtown, 590 West Peachtree Street NW, Atlanta, Georgia 30308 USA. ACM Press, New York, NY, USA. ISBN 978-1-60558-131-6. Part of GECCO 2008, see also [550] and <http://www.eecs.harvard.edu/~sevan/ecomass08/> [accessed 2008-07-21].
- [491] Bogdan Filipič and Jurij Šilc, editors. *Proceedings of the Second International Conference on Bioinspired Optimization Methods and their Application, BIOMA 2006*,

- October 9–10, 2006, Jožef Stefan International Postgraduate School, Ljubljana, Slovenia. Jožef Stefan Institute. ISBN 978-961-6303-81-1.
- [492] Bogdan Filipič and Jurij Šilc, editors. *Proceedings of the International Conference on Bioinspired Optimization Methods and their Applications (BIOMA 2004)*, October 11–12, 2004, Jožef Stefan International Postgraduate School, Ljubljana, Slovenia. Jožef Stefan Institute.
- [493] Zbigniew Michalewicz and Robert G. Reynolds, editors. *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2008*, June 1–6, 2007, Hong Kong Convention and Exhibition Centre, Hong Kong, China. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 978-1-4244-1823-7. See <http://www.wcci2008.org/> [accessed 2008-07-24].
- [494] *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2007*, September 25–28, 2007, Swissôtel The Stamford, Singapore. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. See <http://cec2007.nus.edu.sg/> [accessed 2007-08-06].
- [495] Gary G. Yen, Simon M. Lucas, Gary Fogel, Graham Kendall, Ralf Salomon, Byoung-Tak Zhang, Carlos A. Coello Coello, and Thomas Philip Runarsson, editors. *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2000*, July 16–21, 2000, Sheraton Vancouver Wall Centre Hotel, Vancouver, BC, Canada. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-9487-9.
- [496] David Corne, Zbigniew Michalewicz, Bob McKay, Gusz Eiben, David Fogel, Carlos Fonseca, Garrison Greenwood, Gunther Raidl, Kay Chen Tan, and Ali Zalzal, editors. *Proceedings of the IEEE Congress on Evolutionary Computation, CEC2005*, September 2–5, 2005, Edinburgh, Scotland, UK. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-9363-5.
- [497] *Proceedings of the IEEE Congress on Evolutionary Computation, CEC2004*, June 20–23, 2004, Portland, Oregon. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-8515-2. CEC 2004 – A joint meeting of the IEEE, the EPS, and the IEE.
- [498] Ruhul Sarker, Robert G. Reynolds, Hussein Abbass, Kay Chen Tan, Bob McKay, Daryl Essam, and Tom Gedeon, editors. *Proceedings of the IEEE Congress on Evolutionary Computation, CEC2003*, December 8–12, 2003, Canberra, Australia. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-7804-0. CEC 2003 – A joint meeting of the IEEE, the IEAust, the EPS, and the IEE.
- [499] David B. Fogel, Mohamed A. El-Sharkawi, Xin Yao, Garry Greenwood, Hitoshi Iba, Paul Marrow, and Mark Shackleton, editors. *Proceedings of the IEEE Congress on Evolutionary Computation, CEC2002*, May 12–17, 2002, Honolulu, HI, USA. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-7278-6. CEC 2002 – A joint meeting of the IEEE, the Evolutionary Programming Society, and the IEE. Held in connection with the World Congress on Computational Intelligence (WCCI 2002).
- [500] *Proceedings of the IEEE Congress on Evolutionary Computation, CEC2001*, May 27–30, 2001, COEX, World Trade Center, 159 Samseong-dong, Gangnam-gu, Seoul, Korea. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-6658-1. CEC-2001 – A joint meeting of the IEEE, Evolutionary Programming Society, Galesia, and the IEE. IEEE Catalog Number: 01TH8546C.
- [501] *Proceedings of the IEEE Congress on Evolutionary Computation, CEC00*, July 6–9, 2000, La Jolla Marriott Hotel, La Jolla, California, USA. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-6375-2. doi:10.1109/ICASSP.2000.862060. CEC-2000 – A joint meeting of the IEEE,

- Evolutionary Programming Society, Galesia, and the IEE. IEEE Catalog Number: 00TH8512, Library of Congress Number: 00-018644.
- [502] Peter John Angeline, Zbyszek Michalewicz, Marc Schoenauer, Xin Yao, and Ali Zalzala, editors. *Proceedings of the IEEE Congress on Evolutionary Computation, CEC99*, volume 1-3, July 6–9, 1999, Mayflower Hotel, Washington D.C., USA. IEEE Press, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA. ISBN 0-7803-5536-9, 0-7803-5537-7. CEC-99 – A joint meeting of the IEEE, Evolutionary Programming Society, Galesia, and the IEE. Library of Congress Number: 99-61143.
- [503] *The 1998 IEEE International Conference on Evolutionary Computation, Proceedings of IEEE World Congress on Computational Intelligence, CEC98*, May 4–9, 1998, Anchorage, Alaska, USA. IEEE Press. See <http://ieeexplore.ieee.org/servlet/opac?punumber=5621> [accessed 2007-09-06].
- [504] Thomas Bäck, Zbigniew Michalewicz, and X. Yao, editors. *IEEE International Conference on Evolutionary Computation, CEC97*, April 13–16, 1997, Indianapolis, IN, USA. IEEE Press, Piscataway, NJ. ISBN 0-7803-3949-5. See <http://ieeexplore.ieee.org/servlet/opac?punumber=4619> [accessed 2007-09-06].
- [505] *Proceedings of IEEE International Conference on Evolutionary Computation, CEC96*, May 20–22, 1996, Nagoya, Japan. IEEE Press, Piscataway, NJ. See <http://ieeexplore.ieee.org/servlet/opac?punumber=3838> [accessed 2007-09-06].
- [506] *IEEE International Conference on Evolutionary Computation, ICEC'95*, volume 1-2, November 29–December 1, 1995, Perth, Australia. IEEE Press. See <http://ieeexplore.ieee.org/servlet/opac?punumber=3507> [accessed 2007-09-06]. CEC-95 Editors not given by IEEE, Organisers David Fogel and Chris deSilva.
- [507] Zbigniew Michalewicz, J. David Schaffer, Hans-Paul Schwefel, David B. Fogel, and H. Kitano, editors. *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence*, June 27–29, 1994, Orlando, Florida, USA. IEEE Press, Piscataway, New Jersey. ISBN 0-7803-1899-4. See <http://ieeexplore.ieee.org/servlet/opac?punumber=1125> [accessed 2007-09-06].
- [508] Jürgen Branke, Kalyanmoy Deb, Kaisa Miettinen, and Roman Slowinski, editors. *Practical Approaches to Multi-Objective Optimization*, number 06501 in Dagstuhl Seminar Proceedings, December 10–15, 2006, Dagstuhl, Germany. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany IBFI. Published in 2007. Online available at <http://drops.dagstuhl.de/portals/index.php?semnr=06501> and <http://www.dagstuhl.de/de/programm/kalender/semhp/?semnr=06501> [accessed 2007-09-19].
- [509] Jürgen Branke, Kalyanmoy Deb, Kaisa Miettinen, and Ralph E. Steuer, editors. *Practical Approaches to Multi-Objective Optimization*, number 04461 in Dagstuhl Seminar Proceedings, November 7–12, 2004, Dagstuhl, Germany. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany IBFI. Published in 2005. Online available at <http://drops.dagstuhl.de/portals/index.php?semnr=04461> and <http://www.dagstuhl.de/de/programm/kalender/semhp/?semnr=04461> [accessed 2007-09-19].
- [510] Nicolas Monmarché, El-Ghazali Talbi, Pierre Collet, Marc Schoenauer, and Evelyne Lutton, editors. *Revised Selected Papers of the 8th International Conference on Artificial Evolution, Evolution Artificielle, EA 2007*, Lecture Notes in Computer Science (LNCS), October 29–31, 2007, Tours, France. Springer Berlin/Heidelberg/New York. ISBN 978-3-540-79304-5, 3-540-79304-6. doi:10.1007/978-3-540-79305-2. See <http://ea07.hant.li.univ-tours.fr/> [accessed 2007-09-09]. Published in 2008.
- [511] El-Ghazali Talbi, Pierre Liardet, Pierre Collet, Evelyne Lutton, and Marc Schoenauer, editors. *Revised Selected Papers of the 7th International Conference on Artificial Evolution, Evolution Artificielle, EA 2005*, volume 3871 of *Lecture Notes in Computer Science (LNCS)*, October 26–28, 2005, Lille, France. Springer Berlin/Heidelberg. ISBN 3-540-33589-7. Published in 2006.

- [512] Pierre Liardet, Pierre Collet, Cyril Fonlupt, Evelyne Lutton, and Marc Schoenauer, editors. *Proceedings of the 6th International Conference on Artificial Evolution, Evolution Artificielle, EA 2003*, volume 2936 of *Lecture Notes in Computer Science (LNCS)*, October 27–30, 2003, Marseilles, France. Springer Berlin/Heidelberg. ISBN 3-540-21523-9. Published in 2003.
- [513] Pierre Collet, Cyril Fonlupt, Jin-Kao Hao, Evelyne Lutton, and Marc Schoenauer, editors. *Selected Papers of the 5th International Conference on Artificial Evolution, Evolution Artificielle, EA 2001*, volume 2310 of *Lecture Notes in Computer Science (LNCS)*, October 29–31, 2001, Le Creusot, France. Springer Berlin/Heidelberg. ISBN 3-540-43544-1. Published in 2002.
- [514] Cyril Fonlupt, Jin-Kao Hao, Evelyne Lutton, Edmund M. A. Ronald, and Marc Schoenauer, editors. *Selected Papers of the 4th European Conference on Artificial Evolution, AE'99*, volume 1829 of *Lecture Notes in Computer Science (LNCS)*, November 3–5, 1999, Dunkerque, France. Springer Berlin/Heidelberg. ISBN 3-540-67846-8. Published in 2000.
- [515] Jin-Kao Hao, Evelyne Lutton, Edmund M. A. Ronald, Marc Schoenauer, and Dominique Snyers, editors. *Selected Papers of the Third European Conference on Artificial Evolution, AE'97*, volume 1363 of *Lecture Notes in Computer Science (LNCS)*, October 22–24, 1997, Nîmes, France. Springer Berlin/Heidelberg. ISBN 3-540-64169-6. Published in May 1998.
- [516] Jean-Marc Alliot, Evelyne Lutton, Edmund M. A. Ronald, Marc Schoenauer, and Dominique Snyers, editors. *Selected Papers of the European Conference on Artificial Evolution, AE 95*, volume 1063 of *Lecture Notes in Computer Science (LNCS)*, September 4–6, 1995, Brest, France. Springer Berlin/Heidelberg. ISBN 3-540-61108-8. Published in 1996.
- [517] Jean-Marc Alliot, Evelyne Lutton, and Marc Schoenauer, editors. *European Conference on Artificial Evolution, AE 94*, 1994, Toulouse, France. Cepadues. ISBN 9782854284119. Published in January 1995.
- [518] Shigeru Obayashi, Kalyanmoy Deb, Carlo Poloni, Tomoyuki Hiroyasu, and Tadahiko Murata, editors. *Proceedings of the Fourth International Conference on Evolutionary Multi-Criterion Optimization (EMO'2007)*, volume 4403/2007 of *Lecture Notes in Computer Science (LNCS)*, March 5–8, 2007, Matsushima/Sendai, Japan. Springer-Verlag, Berlin. ISBN 978-3-540-70927-5. See <http://www.is.doshisha.ac.jp/emo2007/> [accessed 2007-09-11].
- [519] Carlos A Coello Coello, Arturo Hernández Aguirre, and Eckart Zitzler, editors. *Proceedings of the Third International Conference on Evolutionary Multi-Criterion Optimization (EMO'05)*, volume 3410 of *Lecture Notes in Computer Science (LNCS)*, March 9–11, 2005, Guanajuato, México. Springer-Verlag, Berlin. ISBN 978-3-540-24983-2.
- [520] Carlos M. Fonseca, Peter J. Fleming, Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele, editors. *Proceedings of the Second International Conference on Evolutionary Multi-Criterion Optimization (EMO'03)*, volume 2632/2003 of *Lecture Notes in Computer Science (LNCS)*, April 8–11, 2003, Faro, Portugal. Springer-Verlag, Berlin. ISBN 978-3-540-01869-8.
- [521] Eckart Zitzler, Kalyanmoy Deb, Lothar Thiele, Carlos A. Coello Coello, and David Corne, editors. *Evolutionary Multi-Criterion Optimization, Proceedings of the First International Conference on Evolutionary Multi-Criterion Optimization (EMO'01)*, volume 1993/2001 of *Lecture Notes in Computer Science (LNCS)*, March 7–9, 2001, Zurich, Switzerland. Springer-Verlag, Berlin. ISBN 3-540-41745-1. See <http://www.tik.ee.ethz.ch/emo/> [accessed 2007-09-11].
- [522] *Proceedings of the Seventh Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control with Applications to Industrial and Societal*

- Problems, EUROGEN2007*, June 11–13, 2007, University of Jyväskylä, Jyväskylä, Finland. See <http://www.mit.jyu.fi/scoma/Eurogen2007/> [accessed 2007-09-16].
- [523] R. Schilling, W. Haase, J. Périaux, and H. Baier, editors. *Proceedings of the Sixth Conference on Evolutionary and Deterministic Methods for Design, Optimization and Control with Applications to Industrial and Societal Problems, EUROGEN2005*, September 12–14, 2005, TU München, Munich, Germany. TU München. in association with ECCOMAS and ERCOFTAC. See <http://www.flm.mw.tum.de/EUROGEN05/> [accessed 2007-09-16].
- [524] *Proceedings of the EUROGEN2003 Conference: Evolutionary Methods for Design Optimization and Control with Applications to Industrial Problems*, September 15–17, 2003, Barcelona, Spain. ISBN 84-95999-33-1. Published on CD. See <http://congress.cimne.upc.es/eurogen03/> [accessed 2007-09-16].
- [525] K. C. Giannakoglou, D. T. Tsahalis, J. Périaux, K. D. Papailiou, and T. Fogarty, editors. *Proceedings of the EUROGEN2001 Conference: Evolutionary Methods for Design Optimization and Control with Applications to Industrial Problems*, September 19–21, 2001, Athens, Greece. International Center for Numerical Methods in Engineering (Cmine), Gran Capitan s/n, 08034 Barcelona, Spain. ISBN 84-89925-97-6. Published in 2002. See <http://www.mech.ntua.gr/~eurogen2001> [accessed 2007-09-16].
- [526] Kaisa Miettinen, Marko M. Mäkelä, Pekka Neittaanmäki, and Jacques Périaux, editors. *European Short Course on Genetic Algorithms and Evolution Strategies, Proceedings of EUROGEN 1999*, May 30–June 3, 1999, University of Jyväskylä, Jyväskylä, Finland.
- [527] D. Quagliarella, J. Périaux, C. Poloni, and G. Winter, editors. *Genetic Algorithms and Evolution Strategy in Engineering and Computer Science: Recent Advances and Industrial Applications – Proceedings of the 2nd European Short Course on Genetic Algorithms and Evolution Strategies, EUROGEN 1997*, Recent Advances in Industrial Applications, November 28–December 4, 1997, Triest, Italy. Wiley & Sons. ISBN 978-0471977100.
- [528] *Genetic Algorithms in Engineering and Computer Science, European Short Course on Genetic Algorithms and Evolution Strategies, EUROGEN 1995*, 1995, Las Palmas de Gran Canaria, Spain. John Wiley & Sons.
- [529] Carlos Cotta and Jano I. van Hemert, editors. *Proceedings of the 7th European Conference on Evolutionary Computation in Combinatorial Optimization, EvoCOP 2007*, volume 4446/2007 of *Lecture Notes in Computer Science (LNCS)*, April 11–13, 2007, Valencia, Spain. Springer. ISBN 978-3-540-71614-3.
- [530] Jens Gottlieb and Günther R. Raidl, editors. *Proceedings of the 6th European Conference on Evolutionary Computation in Combinatorial Optimization, EvoCOP 2006*, volume 3906/2006 of *Lecture Notes in Computer Science (LNCS)*, April 10–12, 2006, Budapest, Hungary. Springer. ISBN 3-540-33178-6.
- [531] Günther R. Raidl and Jens Gottlieb, editors. *Proceedings of the 5th European Conference on Evolutionary Computation in Combinatorial Optimization, EvoCOP 2005*, volume 3448/2005 of *Lecture Notes in Computer Science (LNCS)*, March 30–April 1, 2005, Lausanne, Switzerland. Springer. ISBN 3-540-25337-8.
- [532] Jens Gottlieb and Günther R. Raidl, editors. *Proceedings of the 4th European Conference on Evolutionary Computation in Combinatorial Optimization, EvoCOP 2004*, volume 3004/2004 of *Lecture Notes in Computer Science (LNCS)*, April 5–7, 2004, Coimbra, Portugal. Springer. ISBN 3-540-21367-8.
- [533] Günther R. Raidl, Jean-Arcady Meyer, Martin Middendorf, Stefano Cagnoni, Juan J. Romero Cardalda, David Corne, Jens Gottlieb, Agnès Guillot, Emma Hart, Colin G. Johnson, and Elena Marchiori, editors. *Proceedings of Applications of Evolutionary Computing, EvoWorkshop 2003: EvoBIO, EvoCOP, EvoIASP, EvoMUSART, EvoROB, and EvoSTIM*, volume 2611/2003 of *Lecture Notes in Computer*

- Science (LNCS)*, April 14–16, 2003, University of Essex, Essex, UK. Springer. ISBN 3-540-00976-0.
- [534] Stefano Cagnoni, Jens Gottlieb, Emma Hart, Martin Middendorf, and Günther R. Raidl, editors. *Proceedings of Applications of Evolutionary Computing, EvoWorkshops 2002: EvoCOP, EvoIASP, EvoSTIM/EvoPLAN*, volume 2279/2002 of *Lecture Notes in Computer Science (LNCS)*, April 3–4, 2002, Kinsale, Ireland. Springer. ISBN 3-540-43432-1.
 - [535] Egbert J. W. Boers, Jens Gottlieb, Pier Luca Lanzi, Robert E. Smith, Stefano Cagnoni, Emma Hart, Günther R. Raidl, and H. Tijink, editors. *Proceedings of Applications of Evolutionary Computing, EvoWorkshops 2001: EvoCOP, EvoFlight, EvoIASP, EvoLearn, and EvoSTIM*, volume 2037/2001 of *Lecture Notes in Computer Science (LNCS)*, April 18–20, 2001, Lake Como, Milan, Italy. Springer. ISBN 3-540-41920-9.
 - [536] Mario Giacobini, Anthony Brabazon, Stefano Cagnoni, Gianni A. Di Caro, Rolf Drechsler, Anikó Ekárt, Anna I. Esparcia-Alcázar, Muddassar Farooq, Andreas Fink, Jon McCormack, Michael O'Neill, Juan Romero, Franz Rothlauf, Giovanni Squillero, A. Şima Uyar, and Shengxiang Yang, editors. *Applications of Evolutionary Computing – Proceedings of EvoWorkshops 2008: EvoCOMNET, EvoFIN, EvoHOT, EvoIASP, EvoMUSART, EvoNUM, EvoSTOC, and EvoTransLog*, volume 4974/2008 of *Lecture Notes in Computer Science (LNCS)*, March 26–28, 2008, Naples, Italy. Springer Berlin / Heidelberg. ISBN 3-540-78760-7, 978-3-540-78760-0. doi:10.1007/978-3-540-78761-7.
 - [537] Mario Giacobini, Anthony Brabazon, Stefano Cagnoni, Gianni Di Caro, Rolf Drechsler, Muddassar Farooq, Andreas Fink, Evelyne Lutton, Penousal Machado, Stefan Minner, Michael O'Neill, Juan Romero, Franz Rothlauf, Giovanni Squillero, Hideyuki Takagi, Sima Uyar, and Shengxiang Yang, editors. *Applications of Evolutionary Computing, EvoWorkshops 2007: EvoCoMnet, EvoFIN, EvoIASP, EvoINTERACTION, EvoMUSART, EvoSTOC and EvoTransLog*, volume 4448/2007 of *Lecture Notes in Computer Science (LNCS)*, April 11–13, 2007, Valencia, Spain. Springer. ISBN 978-3-540-71804-8.
 - [538] Franz Rothlauf, Jürgen Branke, Stefano Cagnoni, Ernesto Costa, Carlos Cotta, Rolf Drechsler, Evelyne Lutton, Penousal Machado, Jason H. Moore, Juan Romero, George D. Smith, Giovanni Squillero, and Hideyuki Takagi, editors. *Proceedings of Applications of Evolutionary Computing, EvoWorkshops 2006: EvoBIO, EvoCOMNET, EvoHOT, EvoIASP, EvoINTERACTION, EvoMUSART, and EvoSTOC*, volume 3907/2006 of *Lecture Notes in Computer Science (LNCS)*, April 10–12, 2006, Budapest, Hungary. Springer. ISBN 3-540-33237-5.
 - [539] Franz Rothlauf, Jürgen Branke, Stefano Cagnoni, David W. Corne, Rolf Drechsler, Yaochu Jin, Penousal Machado, Elena Marchiori, Juan Romero, George D. Smith, and Giovanni Squillero, editors. *Proceedings of Applications of Evolutionary Computing, EvoWorkshops 2005: EvoBIO, EvoCOMNET, EvoHOT, EvoIASP, EvoMUSART, and EvoSTOC*, volume 3449/2005 of *Lecture Notes in Computer Science (LNCS)*, March 30–April 1, 2005, Lausanne, Switzerland. Springer. ISBN 3-540-25396-3, 978-3-540-25396-9.
 - [540] Günther R. Raidl, Stefano Cagnoni, Jürgen Branke, David Corne, Rolf Drechsler, Yaochu Jin, Colin G. Johnson, Penousal Machado, Elena Marchiori, Franz Rothlauf, George D. Smith, and Giovanni Squillero, editors. *Proceedings of Applications of Evolutionary Computing, EvoWorkshops 2004: EvoBIO, EvoCOMNET, EvoHOT, EvoIASP, EvoMUSART, and EvoSTOC*, volume 3005/2004 of *Lecture Notes in Computer Science (LNCS)*, April 5–7, 2004, Coimbra, Portugal. Springer. ISBN 3-540-21378-3.
 - [541] Stefano Cagnoni, Riccardo Poli, Yun Li, George D. Smith, David Corne, Martin J. Oates, Emma Hart, Pier Luca Lanzi, Egbert J. W. Boers, Ben Paechter,

- and Terence C. Fogarty, editors. *Real-World Applications of Evolutionary Computing, Proceedings of EvoWorkshops 2000: EvoIASP, EvoSCONDI, EvoTel, EvoSTIM, EvoROB, and EvoFlight*, volume 1803/200 of *Lecture Notes in Computer Science (LNCS)*, April 17, 2000, Edinburgh, Scotland, UK. Springer. ISBN 3-540-67353-9, 978-3-540-67353-8. doi:10.1007/3-540-45561-2.
- [542] Riccardo Poli, Hans-Michael Voigt, Stefano Cagnoni, David Corne, George D. Smith, and Terence C. Fogarty, editors. *Proceedings of the First European Workshops on Evolutionary Image Analysis, Signal Processing and Telecommunications,, EvoIASP'99 and EuroEcTel'99*, volume 1596/1999 of *Lecture Notes in Computer Science (LNCS)*, May 26–27, 1999, Göteborg, Sweden. Springer. ISBN 3-540-65837-8.
- [543] Phil Husbands and Jean-Arcady Meyer, editors. *Proceedings of the First European Workshop on Evolutionary Robotics, EvoRbot98*, volume 1468/1998 of *Lecture Notes in Computer Science (LNCS)*, April 16–17, 1998, Paris, France. Springer. ISBN 3-540-64957-3.
- [544] *Proceedings of the Eighth Joint Conference on Information Science (JCIS 2005), Section: The Sixth International Workshop on Frontiers in Evolutionary Algorithms (FEA 2005)*, July 21–16, 2005, Salt Lake City Marriott City Center, Salt Lake City, Utah, USA. Workshop held in conjunction with Eighth Joint Conference on Information Sciences. See <http://fs.mis.kuas.edu.tw/~cobo1/JCIS2005/jcis05/Track4.html> [accessed 2007-09-16].
- [545] Ken Chen et al., editor. *Proceedings of the Seventh Joint Conference on Information Science (JCIS 2003), Section: The Fifth International Workshop on Frontiers in Evolutionary Algorithms (FEA 2003)*, September 26–30, 2003, Embassy Suites Hotel and Conference Center, Cary, North Carolina, USA. Workshop held in conjunction with Seventh Joint Conference on Information Sciences.
- [546] H. John Caulfield, Shu-Heng Chen, Heng-Da Cheng, Richard J. Duro, Vasant Honavar, Etienne E. Kerre, Mi Lu, Manuel Grana Romy, Timothy K. Shih, Dan Ventura, Paul P. Wang, and Yuanyuan Yang, editors. *Proceedings of the Sixth Joint Conference on Information Science (JCIS 2002), Section: The Fourth International Workshop on Frontiers in Evolutionary Algorithms (FEA 2002)*, March 8–13, 2002, Research Triangle Park, North Carolina, USA. JCIS / Association for Intelligent Machinery, Inc. ISBN 0-9707890-1-7. Workshop held in conjunction with Sixth Joint Conference on Information Sciences.
- [547] Paul P. Wang, editor. *Proceedings of the Fifth Joint Conference on Information Science (JCIS 2000), Section: The Third International Workshop on Frontiers in Evolutionary Algorithms (FEA 2000)*, February 27–March 3, 2000, Atlantic City, NJ, USA. The Association for Intelligent Machinery. Workshop held in conjunction with Fifth Joint Conference on Information Sciences.
- [548] *Proceedings of the Fourth Joint Conference on Information Science (JCIS 1998), Section: The Second International Workshop on Frontiers in Evolutionary Algorithms (FEA 1998)*, volume 2, October 24–28, 1998, Research Triangle Park, North Carolina, USA. The Association for Intelligent Machinery. Workshop held in conjunction with Fourth Joint Conference on Information Sciences.
- [549] *Proceedings of the Third Joint Conference on Information Science (JCIS 1997), Section: The First International Workshop on Frontiers in Evolutionary Algorithms (FEA 1998)*, March 1–5, 1997, Sheraton Imperial Hotel & Convention Center, Research Triangle Park, North Carolina, USA. Workshop held in conjunction with Third Joint Conference on Information Sciences.
- [550] Maarten Keijzer, Giuliano Antoniol, Clare Bates Congdon, Kalyanmoy Deb, Benjamin Doerr, Nikolaus Hansen, John H. Holmes, Gregory S. Hornby, Daniel Howard, James Kennedy, Sanjeev Kumar, Fernando G. Lobo, Julian Francis Miller, Jason Moore, Frank Neumann, Martin Pelikan, Jordan Pollack, Kumara Sastry, Kenneth Stanley, Adrian Stoica, El-Ghazali Talbi, and Ingo Wegener, editors. *Proceedings of*

- the Genetic and Evolutionary Computation Conference, GECCO 2008*, July 12–16, 2008, Renaissance Atlanta Hotel Downtown, 590 West Peachtree Street NW, Atlanta, Georgia 30308 USA. ACM Press, New York, NY, 10286-1405, USA. ISBN 978-1-60558-131-6. ACM Order Number: 910081. See <http://www.sigevo.org/gecco-2008/> [accessed 2008-07-16]. See also [461, 489, 488, 490].
- [551] Dirk Thierens, Hans-Georg Beyer, Josh Bongard, Jurgen Branke, John Andrew Clark, Dave Cliff, Clare Bates Congdon, Kalyanmoy Deb, Benjamin Doerr, Tim Kovacs, Sanjeev Kumar, Julian Francis Miller, Jason Moore, Frank Neumann, Martin Pelikan, Riccardo Poli, Kumara Sastry, Kenneth Owen Stanley, Thomas Stutzle, Richard A Watson, and Ingo Wegener, editors. *Proceedings of Genetic and Evolutionary Computation Conference, GECCO 2007*, July 7–12, 2007, University College London, Gower Street, London WC 1E 6BT, London, UK. ACM Press, New York, NY, USA. ISBN 978-1-59593-697-4. See also [552].
 - [552] Dirk Thierens, Hans-Georg Beyer, Josh Bongard, Jurgen Branke, John Andrew Clark, Dave Cliff, Clare Bates Congdon, Kalyanmoy Deb, Benjamin Doerr, Tim Kovacs, Sanjeev Kumar, Julian Francis Miller, Jason Moore, Frank Neumann, Martin Pelikan, Riccardo Poli, Kumara Sastry, Kenneth Owen Stanley, Thomas Stutzle, Richard A Watson, and Ingo Wegener, editors. *Genetic and Evolutionary Computation Conference - Companion Material, GECCO 2007*, July 7–12, 2007, University College London, Gower Street, London WC 1E 6BT, London, UK. ACM Press, New York, NY, USA. ISBN 978-1-59593-698-1. See also [551], ACM Order Number 910071.
 - [553] Mike Cattolico, editor. *GECCO'06: Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, July 8–12, 2006, Renaissance Seattle Hotel, 515 Madison Street, Seattle, Washington 98104, USA. ACM Press, New York, NY, USA. ISBN 1-59593-186-4. ACM order number 910060.
 - [554] Hans-Georg Beyer, Una-May O'Reilly, Dirk V. Arnold, Wolfgang Banzhaf, Christian Blum, Eric W. Bonabeau, Erick Cantú-Paz, Dipankar Dasgupta, Kalyanmoy Deb, James A. Foster, Edwin D. de Jong, Hod Lipson, Xavier Llorà, Spiros Mancoridis, Martin Pelikan, Guenther R. Raidl, Terence Soule, Andy M. Tyrrell, Jean-Paul Watson, and Eckart Zitzler, editors. *Proceedings of Genetic and Evolutionary Computation Conference, GECCO 2005*, June 25–29, 2005, Loews L'Enfant Plaza Hotel, 480 L'enfant Plaza Sw, Washington, D.C. 20024, USA. ACM Press, Washington DC, USA. ISBN 1-59593-010-8. Order Number: 910050. GECCO-2005 A joint meeting of the fourteenth international conference on genetic algorithms (ICGA-2005) and the tenth annual genetic programming conference (GP-2005). ACM Order Number 910052. See also [555, 556, 557].
 - [555] Hans-Georg Beyer and Una-May O'Reilly, editors. *Workshop Proceedings of Genetic and Evolutionary Computation Conference, GECCO 2005*, June 25–29, 2005, Loews L'Enfant Plaza Hotel, 480 L'enfant Plaza Sw, Washington, D.C. 20024, USA. ACM. ACM Order Number 910052. See also [554, 556, 557].
 - [556] Franz Rothlauf, editor. *Late Breaking Papers at Genetic and Evolutionary Computation Conference, GECCO 2005*, June 25–29, 2005, Loews L'Enfant Plaza Hotel, 480 L'enfant Plaza Sw, Washington, D.C. 20024, USA. See also [554, 555, 557]. Also distributed on CD-ROM at GECCO-2005.
 - [557] Franz Rothlauf, Misty Blowers, Jürgen Branke, Stefano Cagnoni, Ivan I. Garibay, Ozlem Garibay, Jörn Grahl, Gregory Hornby, Edwin D. de Jong, Tim Kovacs, Sanjeev Kumar, Claudio F. Lima, Xavier Llorà, Fernando Lobo, Laurence D. Merkle, Julian Francis Miller, Jason H. Moore, Michael O'Neill, Martin Pelikan, Terry P. Ripopka, Marylyn D. Ritchie, Kumara Sastry, Stephen L. Smith, Hal Stringer, Keiki Takadama, Marc Toussaint, Stephen C. Upton, Alden H. Wright, and Shengxiang Yang, editors. *Proceedings of the 2005 workshops on Genetic and evolutionary computation*, June 25–26, 2005, Loews L'Enfant Plaza Hotel, 480 L'enfant Plaza Sw, Washington, D.C. 20024, USA. ACM, New York, NY, USA. See also [554, 555, 556].

- [558] Kalyanmoy Deb, Riccardo Poli, Wolfgang Banzhaf, Hans-Georg Beyer, Edmund K. Burke, Paul J. Darwen, Dipankar Dasgupta, Dario Floreano, James A. Foster, Mark Harman, Owen Holland, Pier Luca Lanzi, Lee Spector, Andrea Tettamanzi, Dirk Thierens, and Andrew M. Tyrrell, editors. *Proceedings of Genetic and Evolutionary Computation - GECCO 2004, Genetic and Evolutionary Computation Conference, Part I*, volume 3102/2004 of *Lecture Notes in Computer Science (LNCS)*, June 26–30, 2004, Red Lion Hotel, 1415 5th Avenue, Seattle, WA 98101, USA. Springer Berlin/Heidelberg. ISBN 978-3-540-22344-3. See also [559, 560, 1008].
- [559] Kalyanmoy Deb, Riccardo Poli, Wolfgang Banzhaf, Hans-Georg Beyer, Edmund K. Burke, Paul J. Darwen, Dipankar Dasgupta, Dario Floreano, James A. Foster, Mark Harman, Owen Holland, Pier Luca Lanzi, Lee Spector, Andrea Tettamanzi, Dirk Thierens, and Andrew M. Tyrrell, editors. *Proceedings of Genetic and Evolutionary Computation - GECCO 2004, Genetic and Evolutionary Computation Conference, Part II*, volume 3103/2004 of *Lecture Notes in Computer Science (LNCS)*, June 26–30, 2004, Red Lion Hotel, 1415 5th Avenue, Seattle, WA 98101, USA. Springer Berlin/Heidelberg. ISBN 978-3-540-22343-6. See also [558, 560, 1008].
- [560] Maarten Keijzer, editor. *Late Breaking Papers at the 2004 Genetic and Evolutionary Computation Conference*, June 26–30, 2004, Red Lion Hotel, 1415 5th Avenue, Seattle, WA 98101, USA. Springer Berlin/Heidelberg. Distributed on CD-ROM and (online) in [559].
- [561] Erick Cantú-Paz, James A. Foster, Kalyanmoy Deb, Lawrence Davis, Rajkumar Roy, Una-May O'Reilly, Hans-Georg Beyer, Russell K. Standish, Graham Kendall, Stewart W. Wilson, Mark Harman, Joachim Wegener, Dipankar Dasgupta, Mitchell A. Potter, Alan C. Schultz, Kathryn A. Dowsland, Natasa Jonoska, and Julian Francis Miller, editors. *Proceedings of Genetic and Evolutionary Computation - GECCO 2003, Genetic and Evolutionary Computation Conference, Part I*, volume 2723/2003 of *Lecture Notes in Computer Science (LNCS)*, July 12–16, 2003, The Holiday Inn Chicago – Mart Plaza, 350 N. Orleans St., Chicago, IL 60654, USA. Springer Berlin/Heidelberg. ISBN 978-3-540-40602-0. doi:10.1007/3-540-45105-6. See also [562, 1009].
- [562] Erick Cantú-Paz, James A. Foster, Kalyanmoy Deb, Lawrence Davis, Rajkumar Roy, Una-May O'Reilly, Hans-Georg Beyer, Russell K. Standish, Graham Kendall, Stewart W. Wilson, Mark Harman, Joachim Wegener, Dipankar Dasgupta, Mitchell A. Potter, Alan C. Schultz, Kathryn A. Dowsland, Natasa Jonoska, and Julian Francis Miller, editors. *Proceedings of Genetic and Evolutionary Computation - GECCO 2003, Genetic and Evolutionary Computation Conference, Part II*, volume 2724/2003 of *Lecture Notes in Computer Science (LNCS)*, July 12–16, 2003, The Holiday Inn Chicago – Mart Plaza, 350 N. Orleans St., Chicago, IL 60654, USA. Springer Berlin/Heidelberg. ISBN 978-3-540-40603-7. doi:10.1007/3-540-45110-2. See also [561, 1009].
- [563] William B. Langdon, Erick Cantú-Paz, Keith E. Mathias, Rajkumar Roy, David Davis, Riccardo Poli, Karthik Balakrishnan, Vasant Honavar, Günter Rudolph, Joachim Wegener, Larry Bull, Mitchell A. Potter, Alan C. Schultz, Julian F. Miller, Edmund K. Burke, and Natasa Jonoska, editors. *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, July 9–13, 2002, The Roosevelt Hotel, 45th and Madison Avenue, New York, NY, USA. Morgan Kaufmann Publishers Inc. ISBN 1-55860-878-8. See also [564, 566, 565].
- [564] Erick Cantú-Paz, editor. *Late Breaking papers at the Genetic and Evolutionary Computation Conference (GECCO-2002)*, July 9–13, 2002, The Roosevelt Hotel, 45th and Madison Avenue, New York, NY, USA. AAAI, 445 Burgess Drive, Menlo Park, CA 94025. See also [563, 566, 565].
- [565] Alwyn M. Barry, editor. *GECCO 2002: Proceedings of the Bird of a Feather Workshops, Genetic and Evolutionary Computation Conference*, July 8, 2002, The Roo-

- sevelt Hotel, 45th and Madison Avenue, New York, NY, USA. AAI, 445 Burgess Drive, Menlo Park, CA 94025. Bird-of-a-feather Workshops, GECCO-2002. A joint meeting of the eleventh International Conference on Genetic Algorithms (ICGA-2002) and the seventh Annual Genetic Programming Conference (GP-2002). [564, 563, 566].
- [566] Michael O'Neill and Conor Ryan, editors. *Grammatical Evolution Workshop (GEWS 2002)*, July 9, 2002, New York, NY, USA. Part of GECCO, see [563] and also <http://www.grammatical-evolution.com/gews2002/> [accessed 2007-09-10].
- [567] Sean Luke, Conor Ryan, and Una-May O'Reilly, editors. *GECCO 2002: Graduate Student Workshop*, July 9, 2002, New York. AAI, 445 Burgess Drive, Menlo Park, CA 94025. Part of [565].
- [568] Lee Spector, Erik D. Goodman, Annie Wu, William B. Langdon, Hans-Michael Voigt, Mitsuo Gen, Sandip Sen, Marco Dorigo, Shahram Pezeshk, Max H. Garzon, and Edmund Burke, editors. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'01)*, July 7–11, 2001, Holiday Inn Golden Gateway Hotel, San Francisco, California, USA. Morgan Kaufmann. ISBN 978-1558607743, 1-55860-774-9. See also [569].
- [569] Erik D. Goodman, editor. *Late Breaking Papers at Genetic and Evolutionary Computation Conference (GECCO'01)*, July 7–11, 2001, Holiday Inn Golden Gateway Hotel, San Francisco, California, USA. See also [568].
- [570] L. Darrell Whitley, David Goldberg, Erick Cantú-Paz, Lee Spector, Ian C. Parmee, and Hans-Georg Beyer, editors. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'00)*, July 8–12, 2000, The Riviera Hotel and Casino, Las Vegas, Nevada, USA. Morgan Kaufmann. ISBN 978-1558607088. See also [571].
- [571] L. Darrell Whitley, editor. *Late Breaking Papers at Genetic and Evolutionary Computation Conference (GECCO'00)*, July 8–12, 2000, The Riviera Hotel and Casino, Las Vegas, Nevada, USA. See also [570].
- [572] Wolfgang Banzhaf, Jason M. Daida, Agoston E. Eiben, Max H. Garzon, Vasant Honavar, Mark J. Jakiela, and Robert E. Smith, editors. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)*, July 13–17, 1999, Orlando, Florida, USA. Morgan Kaufmann. ISBN 1-55860-611-4. See also [573, 478].
- [573] *Late Breaking Papers at Genetic and Evolutionary Computation Conference (GECCO 1999)*, July 13–17, 1999, Orlando, Florida, USA. See also [572, 478].
- [574] Bartłomiej Beliczyński, Andrzej Dzieliński, Marcin Iwanowski, and Bernardete Ribeiro, editors. *Proceedings of Adaptive and Natural Computing Algorithms, 8th International Conference, ICANNGA 2007, Part I*, volume 4431/2007 of *Lecture Notes in Computer Science (LNCS)*, April 11–14, 2007, Warsaw University of Technology, Warsaw, Poland. Springer Berlin Heidelberg New York. ISBN 978-3-540-71589-4. doi:10.1007/978-3-540-71618-1. See also [575], <http://icannga07.ee.pw.edu.pl/> [accessed 2007-08-31], and <http://www.springerlink.com/content/978-3-540-71590-0/> [accessed 2007-08-31].
- [575] Bartłomiej Beliczyński, Andrzej Dzieliński, Marcin Iwanowski, and Bernardete Ribeiro, editors. *Proceedings of Adaptive and Natural Computing Algorithms, 8th International Conference, ICANNGA 2007, Part II*, volume 4432/2007 of *Lecture Notes in Computer Science (LNCS)*, April 11–14, 2007, Warsaw University of Technology, Warsaw, Poland. Springer Berlin Heidelberg New York. ISBN 978-3-540-71590-0. doi:10.1007/978-3-540-71629-7. See also [574], <http://icannga07.ee.pw.edu.pl/> [accessed 2007-08-31], and <http://www.springerlink.com/content/978-3-540-71589-4/> [accessed 2007-08-31].
- [576] Bernadete Ribeiro, Rudolf F. Albrecht, Andrej Dobnikar, David W. Pearson, and Nigel C. Steele, editors. *Proceedings of the 7th International Conference on Adaptive and Natural Computing Algorithms, ICANNGA 2005*, March 21–23, 2005, Department of Informatics Engineering, Faculty of Science and Technology, University of

- Coimbra, Coimbra, Portugal. Springer, Wien. ISBN 978-3-211-24934-5 (Print) 978-3-211-27389-0 (Online). see <http://icannga05.dei.uc.pt/> [accessed 2007-08-31].
- [577] David W. Pearson, Nigel C. Steele, and Rudolf F. Albrecht, editors. *Proceedings of the 6th International Conference on Artificial Neural Nets and Genetic Algorithms, ICANNGA*, 2003, Roanne, France. Springer Verlag. ISBN 978-3-211-00743-3.
- [578] Vera Kurkova, Nigel C. Steele, Roman Neruda, and Miroslav Karny, editors. *Proceedings of the 5th International Conference on Artificial Neural Networks and Genetic Algorithms, ICANNGA*, 2001, Prague, Czech Republic. Springer Verlag, Berlin. ISBN 978-3-211-83651-4.
- [579] Andrej Dobnikar, Nigel C. Steele, David W. Pearson, and Rudolf F. Albrecht, editors. *Proceedings of the 4th International Conference on Artificial Neural Nets and Genetic Algorithms, ICANNGA*, 1999, Portoroz, Slovenia. Springer. ISBN 978-3-211-83364-3, 3-211-83364-1.
- [580] George D. Smith, Nigel C. Steele, and Rudolf F. Albrecht, editors. *Proceedings of the International Conference on Artificial Neural Nets and Genetic Algorithms, ICANNGA*, 1997, University of East Anglia, Norwich, England, U.K. Springer, Vienna, Austria. ISBN 978-3-211-83087-1, 3-211-83087-1. Published in 1998.
- [581] David W., Nigel C. Steele, and F. Albrecht, editors. *Proceedings of the 2nd International Conference on Artificial Neural Nets and Genetic Algorithms, ICANNGA*, April 18–21, 1995, École des Mines d'Alès, Alès, France. Springer Verlag. ISBN 3-211-82692-0.
- [582] Rudolf F. Albrecht, Colin R. Reeves, and Nigel C. Steele, editors. *Proceedings of Artificial Neural Nets and Genetic Algorithms, ANNGA*, April 1993, Innsbruck, Austria. Springer, New York. ISBN 3-211-82459-6, 0-387-82459-6.
- [583] *Proceedings of 10th International Conference on Parallel Problem Solving from Nature – PPSN X*, Lecture Notes in Computer Science (LNCS), September 13–17, 2008, Technische Universität Dortmund / Kongresszentrum Westfalenhallen, Dortmund, Germany. Springer. See <http://ls11-www.cs.uni-dortmund.de/ppsn/ppsn10/> [accessed 2008-07-20].
- [584] Thomas Philip Runarsson, Hans-Georg Beyer, Edmund K. Burke, Juan J. Merelo Guervós, L. Darrell Whitley, and Xin Yao, editors. *Proceedings of 9th International Conference on Parallel Problem Solving from Nature – PPSN IX*, volume 4193/2006 of *Lecture Notes in Computer Science (LNCS)*, September 9–13, 2006, University of Iceland, Reykjavik, Iceland. Springer. ISBN 3-540-38990-3. See <http://ppsn2006.raunvis.hi.is/> [accessed 2007-09-05].
- [585] Xin Yao, Edmund K. Burke, José Antonio Lozano, Jim Smith, Juan J. Merelo Guervós, John A. Bullinaria, Jonathan E. Rowe, Peter Tiño, Ata Kabán, and Hans-Paul Schwefel, editors. *Proceedings of the 8th International Conference on Parallel Problem Solving from Nature – PPSN VIII*, volume 3242/2004 of *Lecture Notes in Computer Science (LNCS)*, September 18–22, 2004, Birmingham, UK. Springer. ISBN 3-540-23092-0. See <http://events.cs.bham.ac.uk/ppsn04/> [accessed 2007-09-05].
- [586] Juan J. Merelo Guervós, Panagiotis Adamidis, Hans-Georg Beyer, José Luis Fernández-Villacañas Martín, and Hans-Paul Schwefel, editors. *Proceedings of the 7th International Conference on Parallel Problem Solving from Nature – PPSN VII*, volume 2439/2002 of *Lecture Notes in Computer Science (LNCS)*, September 7–11, 2002, Granada, Spain. Springer. ISBN 3-540-44139-5. See <http://ppsn2002.ugr.es/index.html> [accessed 2007-09-05].
- [587] Marc Schoenauer, Kalyanmoy Deb, Günter Rudolph, Xin Yao, Evelyne Lutton, Juan J. Merelo Guervós, and Hans-Paul Schwefel, editors. *Proceedings of the 6th International Conference on Parallel Problem Solving from Nature – PPSN VI*, volume 1917/2000 of *Lecture Notes in Computer Science (LNCS)*, September 18–20, 2000, Paris, France. Springer. ISBN 3-540-41056-2.

- [588] Agoston E. Eiben, Thomas Bäck, Marc Schoenauer, and Hans-Paul Schwefel, editors. *Proceedings of the 5th International Conference on Parallel Problem Solving from Nature – PPSN V*, volume 1498/1998 of *Lecture Notes in Computer Science (LNCS)*, September 27–30, 1998, Amsterdam, The Netherlands. Springer. ISBN 3-540-65078-4.
- [589] Hans-Michael Voigt, Werner Ebeling, Ingo Rechenberger, and Hans-Paul Schwefel, editors. *Parallel Problem Solving from Nature – PPSN IV, International Conference on Evolutionary Computation. Proceedings of the 4th International Conference on Parallel Problem Solving from Nature*, volume 1141/1996 of *Lecture Notes in Computer Science (LNCS)*, September 22–24, 1996, Berlin, Germany. Springer. ISBN 3-540-61723-X. doi:10.1007/3-540-61723-X. See <http://ls11-www.informatik.uni-dortmund.de/PPSN/ppsn4/ppsn4.html> [accessed 2007-09-05].
- [590] Yuval Davidor, Hans-Paul Schwefel, and Reinhard Männer, editors. *Parallel Problem Solving from Nature – PPSN III, International Conference on Evolutionary Computation. Proceedings of the Third Conference on Parallel Problem Solving from Nature*, volume 866/1994 of *Lecture Notes in Computer Science (LNCS)*, October 9–14, 1994, Jerusalem, Israel. Springer. ISBN 3-540-58484-6. See <http://ls11-www.informatik.uni-dortmund.de/PPSN/ppsn3/ppsn3.html> [accessed 2007-09-05].
- [591] Reinhard Männer and Bernard Manderick, editors. *Proceedings of Parallel Problem Solving from Nature 2, PPSN II*, September 28–30, 1992, Brussels, Belgium. Elsevier. See <http://ls11-www.informatik.uni-dortmund.de/PPSN/ppsn2/ppsn2.html> [accessed 2007-09-05].
- [592] Hans-Paul Schwefel and Reinhard Männer, editors. *Proceedings of the 1st Workshop on Parallel Problem Solving from Nature, PPSN I*, volume 496/1991 of *Lecture Notes in Computer Science (LNCS)*, October 1–3, 1990, FRG, Dortmund, Germany. Springer. ISBN 3-540-54148-9. Published 1991. See <http://ls11-www.informatik.uni-dortmund.de/PPSN/ppsn1/ppsn1.html> [accessed 2007-09-05].
- [593] Hans-Georg Beyer, Eva Brucherseifer, Wilfried Jakob, Hartmut Pohlheim, Bernhard Sendhoff, and Thanh Binh To. *Evolutionary Algorithms – Terms and Definitions*. Universität Dortmund, Informatik XI (Systemanalyse), D-44221 Dortmund, Germany, version e-1.2 edition, February 25, 2002. Online available at <http://ls11-www.cs.uni-dortmund.de/people/beyer/EA-glossary/> [accessed 2008-04-10].
- [594] David B. Fogel, editor. *Evolutionary Computation: The Fossil Record*. Wiley-IEEE Press, May 1, 1998. ISBN 978-0780334816.
- [595] Carlos Artemio Coello Coello and Gary B. Lamont, editors. *Applications of Multi-Objective Evolutionary Algorithms*, volume 1 of *Advances in Natural Computation*. World Scientific Publishing Co., December 2004. ISBN 978-981-256-106-0, 981-256-106-4. Partly online available at <http://books.google.de/books?id=viWm0k9me0cC> [accessed 2008-07-20].
- [596] D. Dumitrescu, Beatrice Lazzerini, Lakhmi C. Jain, and A. Dumitrescu. *Evolutionary Computation*, volume 18 of *International Series on Computational Intelligence*. CRC Press, June 2000. ISBN 9780849305887, 0849305888. Partly online available at <http://books.google.de/books?hl=de&id=MSU9ep79JvUC> [accessed 2008-06-11].
- [597] Thomas Bäck, David B. Fogel, and Zbigniew Michalewicz, editors. *Evolutionary Computation 1: Basic Algorithms and Operators*. Institute of Physics Publishing (IOP), Dirac House, Temple Back, Bristol BS1 6BE, United Kingdom, January 2000. ISBN 9780750306645, 0750306645. Partly online available at <http://books.google.de/books?id=4HMYCq9US78C> [accessed 2008-06-11].
- [598] Thomas Bäck, David B. Fogel, and Zbigniew Michalewicz, editors. *Evolutionary Computation 2: Advanced Algorithms and Operators*. Institute of Physics Publishing (IOP), Dirac House, Temple Back, Bristol BS1 6BE, United Kingdom, November 2000. ISBN 9780750306652, 0750306653.

- [599] Peter J. Bentley, editor. *Evolutionary Design by Computers*. Morgan Kaufmann, San Francisco, USA, June 15, 1999. ISBN 1-55860-605-X. See also <http://www.cs.ucl.ac.uk/staff/P.Bentley/evdes.html> [accessed 2007-10-03].
- [600] Kenneth Alan De Jong. *Evolutionary Computation: A Unified Approach*, volume 4 of *Complex Adaptive Systems: A Bradford Book*. MIT Press, Cambridge, Massachusetts, USA, February 2006. ISBN 9780262041942, 0262041944.
- [601] Ingrid Gerdes, Klaus Klawonn, and Rudolf Kruse. *Evolutionäre Algorithmen*. Friedrich Vieweg & Sohn Verlag / GWV Fachverlage GmbH (Springer Science+Business Media), Wiesbaden, Germany, July 2004. ISBN 3-528-05570-7.
- [602] Volker Nissen. *Einführung in evolutionäre Algorithmen: Optimierung nach dem Vorbild der Evolution*. Vieweg, Braunschweig, 1997. ISBN 3528054999, 978-3528054991.
- [603] Xin Yao, editor. *Evolutionary Computation: Theory and Applications*. World Scientific Publishing Co Pte Ltd, January 28, 1996. ISBN 9810223064, 978-9810223069. Partly online available at <http://books.google.de/books?id=GP7ChxbJOE4C> [accessed 2008-03-25].
- [604] Gwoing Tina Yu, Lawrence Davis, Cem M. Baydar, and Rajkumar Roy, editors. *Evolutionary Computation in Practice*, volume 88/2008 of *Studies in Computational Intelligence*. Springer, January 2008. ISBN 978-3-540-75770-2. doi:10.1007/978-3-540-75771-9. Series editor: Janusz Kacprzyk.
- [605] Wolfgang Banzhaf and Frank H. Eeckman, editors. *Evolution and Biocomputation – Computational Models of Evolution*, volume 899/1995 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, London, UK, kindle edition, April 13 1995, Monterey, California, USA. ISBN 3-540-59046-3, 0-387-59046-3, 978-3-540-59046-0. doi:10.1007/3-540-59046-3. ASIN: B000V1DAP0. Partly online available at <http://books.google.de/books?id=dZf-yrmsSZPkC> [accessed 2008-04-30].
- [606] David B. Fogel. *Evolutionary Computation: Principles and Practice for Signal Processing*. Tutorial Texts in Optical Engineering. Society of Photo-Optical Instrumentation Engineering (SPIE) Press, July 6, 2000. ISBN 0819437255, 978-0819437259. Partly online available at <http://books.google.de/books?id=gEmZ80JVHXcC> [accessed 2008-04-03].
- [607] Daniel Ashlock from the Department of Mathematics and Statistics, University of Guelph, Ontario N1G 2W1, Canada. *Evolutionary Computation for Modeling and Optimization*. Springer, Berlin, Germany, December 1, 2005. ISBN 0387221964, 978-0-387-22196-0. Partly online available at <http://books.google.de/books?id=74Q5WBOaoz8C> [accessed 2008-04-03].
- [608] Oliver Kramer. *Self-Adaptive Heuristics for Evolutionary Computation*, volume 147 of *Studies in Computational Intelligence*. Springer Berlin / Heidelberg, July 2008. ISBN 978-3-540-69280-5. Series editor: Janusz Kacprzyk.
- [609] Fernando G. Lobo, Cláudio F. Lima, and Zbigniew Michalewicz, editors. *Parameter Setting in Evolutionary Algorithms*, volume 54 of *Studies in Computational Intelligence*. Springer-Verlag GmbH, Berlin, Germany, March 1, 2007. ISBN 3540694315, 978-3540694311.
- [610] William M. Spears. *Evolutionary Algorithms – The Role of Mutation and Recombination*. Natural Computing Series. Springer, Berlin, 2000. ISBN 978-3-540-66950-0.
- [611] Agoston E. Eiben and A. Michalewicz, editors. *Evolutionary Computation*. Ohmsha, Tokyo, Japan, 1999. ISBN 4-274-90269-2. This is the book edition of the journal, Fundamenta Informaticae, Volume 35, Nos. 1-4, 1998.
- [612] Yaochu Jin, editor. *Knowledge Incorporation in Evolutionary Computation*, volume 167 of *Studies in Fuzziness and Soft Computing*. Springer, Springer Distribution Center GmbH, Haberstrasse 7, 69126 Heidelberg, Germany, 2004. ISBN 978-3-540-22902-5. Partly online available at http://books.google.de/books?id=kD_FXkAKn9IC [accessed 2008-05-31].

- [613] Crina Grosan, Ajith Abraham, and Hisao Ishibuchi, editors. *Hybrid Evolutionary Algorithms*, volume 75 of *Studies in Computational Intelligence*. Springer Berlin / Heidelberg, 2007. ISBN 978-3-540-73296-9. doi:10.1007/978-3-540-73297-6. Series editor: Janusz Kacprzyk.
- [614] Ajith Abraham, Lakhmi C. Jain, and Robert Goldberg, editors. *Evolutionary Multiobjective Optimization – Theoretical Advances and Applications*. Advanced Information and Knowledge Processing. Springer, Berlin, July 30, 2005. ISBN 978-1-85233-787-2, 1852337877. Partly online available at <http://books.google.de/books?id=Ei7q1YSjiSAC> [accessed 2008-06-07]. Series editors: Lakhmi C. Jain and Xindong Wu.
- [615] Leila Kallel, Bart Naudts, and Alex Rogers, editors. *Theoretical Aspects of Evolutionary Computing*. Natural Computing Series. Springer, Berlin/London, UK, June 15, 2001. ISBN 3-540-67396-2, 978-3540673965. Series editors: G. Rozenberg, Thomas Bäck, Agoston E. Eiben, J.N. Kok, and H.P. Spaink.
- [616] Ashish Ghosh and Shigeyoshi Tsutsui, editors. *Advances in Evolutionary Computing – Theory and Applications*. Natural Computing Series. Springer-Verlag New York, Inc., New York, NY, USA, November 22, 2002. ISBN 3540433309, 978-3-540-43330-9. Series editors: G. Rozenberg, Thomas Bäck, Agoston E. Eiben, J.N. Kok, and H.P. Spaink.
- [617] Ang Yang, Yin Shan, and Lam Thu Bui, editors. *Success in Evolutionary Computation*, volume 92/2008 of *Studies in Computational Intelligence*. Springer Berlin / Heidelberg, January 16, 2008. ISBN 354076285X, 978-3-540-76285-0, 978-3-540-76286-7. doi:10.1007/978-3-540-76286-7. Library of Congress Control Number: 2007939404.
- [618] Haiming Lu. *State-of-the-art Multiobjective Evolutionary Algorithms – Pareto Ranking, Density Estimation and Dynamic Population*. PhD thesis, Faculty of the Graduate College of the Oklahoma State University, Stillwater, Oklahoma, August 2002. Advisor: Gary G. Yen. Publication Number AAT 3080536. Online available at http://www.lania.mx/~ccoello/EM00/thesis_lu.pdf.gz [accessed 2007-08-25].
- [619] Peter A. N. Bosman and Dirk Thierens. The Naive MIDEA: A Baseline Multiobjective EA. In *Proceedings of the Third International Conference on Evolutionary Multi-Criterion Optimization*, pages 428–442, 2005. In proceedings [519]. Online available at http://springerlink.metapress.com/content/m01x9t6l2bwd/?p_o=20 [accessed 2007-08-22].
- [620] Garrison W. Greenwood, Xiaobo Sharon Hu, and Joseph G. D’Ambrosio. Fitness Functions for Multiple Objective Optimization Problems: Combining Preferences with Pareto Rankings. In *Foundations of Genetic Algorithms 4*, pages 437–455, 1996. In proceedings [736].
- [621] Kalyanmoy Deb. Genetic Algorithms for multimodal function optimization. Master’s thesis, The Clearinghouse for Genetic Algorithms, University of Alabama, Tuscaloosa, 1989.
- [622] David E. Goldberg and Jon Richardson. Genetic algorithms with sharing for multimodal function optimization. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 41–49, 1987. In proceedings [750].
- [623] Kalyanmoy Deb and David E. Goldberg. An investigation of niche and species formation in genetic function optimization. In *Proceedings of the third international conference on Genetic algorithms*, pages 42–50, 1989. In proceedings [652].
- [624] Brad L. Miller and Michael J. Shaw from the University of Illinois, Urbana-Champaign, Illinois, USA. Genetic Algorithms with Dynamic Niche Sharing for Multimodal Function Optimization. IlliGAL Report 95010, Department of General Engineering, University of Illinois at Urbana-Champaign, 117 Transportation Building, 104 South Mathews Avenue, Urbana, IL 61801, USA, December 1,

1995. Online available at <http://citeseer.comp.nus.edu.sg/5466.html> [accessed <http://www.illigal.uiuc.edu/pub/papers/IlliGALs/95010.ps.Z>]2008-03-15.
- [625] Kalyanmoy Deb. Genetic Algorithms in Multimodal Function Optimization. Master's thesis, The Clearinghouse for Genetic algorithms, University of Alabama, Tuscaloosa, 1989. TCGA Report No. 89002.
- [626] Jeffrey Horn, Nicholas Nafpliotis, and David E. Goldberg. A Niche Pareto Genetic Algorithm for Multiobjective Optimization. In *Proceedings of the First IEEE Conference on Evolutionary Computation*, volume 1, pages 82–87, 1994. In proceedings [507]. Also: IEEE Symposium on Circuits and Systems, pp. 2264–2267, 1991, Online available at <http://citeseer.ist.psu.edu/horn94nched.html> and <http://www.lania.mx/~ccoello/EM00/horn2.ps.gz> [accessed 2007-08-28].
- [627] Richard W. Hamming. Error-detecting and error-correcting codes. *Bell System Technical Journal*, 29(2):147–169, 1950. Online available at <http://guest.engelschall.com/~sb/hamming/> and http://garfield.library.upenn.edu/classics/classics_h.html [accessed 2007-08-13].
- [628] Kalyanmoy Deb. An introduction to genetic algorithms. *Sadhana*, 24(4-5):293–315, 1999. Online available at <http://www.iitk.ac.in/kangal/papers/sadhana.ps.gz> [accessed 2007-07-28].
- [629] David E. Goldberg, Kalyanmoy Deb, and Jeffrey Horn. Massive Multimodality, Deception, and Genetic Algorithms. In *Parallel Problem Solving from Nature 2*, 1992. In proceedings [591]. Online available at <http://citeseer.ist.psu.edu/goldberg92massive.html> [accessed 2008-03-16].
- [630] Tobias Blickle and Lothar Thiele. A Comparison of Selection Schemes used in Evolutionary Algorithms. *Evolutionary Computation*, 4(4):361–394, 1996. Online available at <http://citeseer.ist.psu.edu/blickle97comparison.html> [accessed 2007-08-24].
- [631] Kumara Sastry and David E. Goldberg from the Illinois Genetic Algorithms Laboratory (IlliGAL), Department of General Engineering, University of Illinois at Urbana-Champaign. Modeling Tournament Selection With Replacement Using Apparent Added Noise. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, page 781, 2001. In proceedings [568]. Also: IlliGAL report 2001014, January 2001, Illinois Genetic Algorithms Laboratory (IlliGAL), Department of General Engineering, University of Illinois at Urbana-Champaign, Online available at <http://citeseer.ist.psu.edu/439887.html> and <http://citeseer.ist.psu.edu/501944.html> [accessed 2007-08-25].
- [632] David E. Goldberg and Kalyanmoy Deb. A Comparative Analysis of Selection Schemes used in Genetic Algorithms. In *Proceedings of Foundations of Genetic Algorithms*, pages 69–93, 1990. In proceedings [188]. Online available at <http://www.cse.unr.edu/~sushil/class/gas/papers/Select.pdf> [accessed 2007-07-29].
- [633] Jinghui Zhong, Xiaomin Hu, Jun Zhang, and Min Gu. Comparison of Performance between Different Selection Strategies on Simple Genetic Algorithms. In *CIMCA'05: Proceedings of the International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce Vol-2 (CIMCA-IAWTIC'06)*, pages 1115–1121, November 28–30, 2005. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2504-0-02.
- [634] Jörg Lässig, Karl Heinz Hoffmann, and Mihaela Enăchescu. Threshold Selecting: Best Possible Probability Distribution for Crossover Selection in Genetic Algorithms. In *Genetic and Evolutionary Computation Conference*, pages 2181–2185, 2008. In proceedings [550].
- [635] A. Brindle. *Genetic algorithms for function optimization*. PhD thesis, University of Alberta, Department of Computer Science, Edmonton, 1981. Technical Report TR81-2.

- [636] Lashon Bernard Booker. *Intelligent behavior as an adaptation to the task environment*. PhD thesis, University of Michigan, Ann Arbor, MI, USA, 1982. Order No. AAI8214966, Technical Report No. 243. Dissertation Abstracts International, 43(2), 469B. (University Microfilms No. 8214966). Online available at <http://hdl.handle.net/2027.42/3746> and <http://hdl.handle.net/2027.42/3747> [accessed 2008-03-19].
- [637] James E. Baker. Reducing Bias and Inefficiency in the Selection Algorithm. In *Proceedings of the second International Conference on Genetic Algorithms*, pages 14–21, 1987. In proceedings [750].
- [638] Tobias Blickle and Lothar Thiele. A Mathematical Analysis of Tournament Selection. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 9–16, 1995. In proceedings [747]. Online available at <http://citeseer.ist.psu.edu/blickle95mathematical.html> [accessed 2007-07-29].
- [639] A. Wetzel. *Evaluation of the Effectiveness of Genetic Algorithms in Combinatorial Optimization*. PhD thesis, University of Pittsburgh, Pittsburgh, PA, 1983. Unpublished manuscript, technical report.
- [640] S. Lee, S. Soak, K. Kim, H. Park, and M. Jeon. Statistical properties analysis of real world tournament selection in genetic algorithms. *Applied Intelligence*, 28(2):195–205, April 2008. ISSN 0924-669X (Print) 1573-7497 (Online). doi:10.1007/s10489-007-0062-2. Online available at <http://www.springerlink.com/content/amr3nq330x13u32t/fulltext.pdf> [accessed 2008-03-20].
- [641] David Andre. The Evolution of Agents that Build Mental Models and Create Simple Plans Using Genetic Programming. In *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 248–255, 1995. ISBN 1-55860-370-0. In proceedings [747].
- [642] Martin V. Butz, Kumara Sastry, and Davi E. Goldberg. Tournament Selection in XCS. IlliGAL Report 2002020, Illinois Genetic Algorithms Laboratory (IlliGAL), Department of Computer Science, Department of General Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA, July 2002. Online available at <http://citeseer.ist.psu.edu/534844.html> and <http://www.illigal.uiuc.edu/pub/papers/IlliGALs/2002020.ps.Z> [accessed 2008-03-20].
- [643] Mezura-Montes Efrén and Carlos Artemio Coello Coello. Using the Evolution Strategies’ Self-Adaptation Mechanism and Tournament Selection for Global Optimization. *Intelligent Engineering Systems through Artificial Neural Networks (ANNIE’2003)*, 13:373–378, November 2003. Theoretical Developments in Computational Intelligence Award, 1st runner up. Online available at <http://www.lania.mx/~emezura/documentos/annie03.pdf> [accessed 2008-03-20].
- [644] Guilherme Bastos Alvarenga and Geraldo Robson Mateus from the Department of Computer Science, UFPA, Brazil. Hierarchical tournament selection genetic algorithm for the vehicle routing problem with time windows. In *4th International Conference on Hybrid Intelligent Systems*, pages 410–415, 2004. doi:10.1109/ICHIS.2004.53. In proceedings [306].
- [645] L. Darell Whitley from the Computer Science Department, Colorado State University, Fort Collins. A Genetic Algorithm Tutorial. *Statistics and Computing*, 4(2):65–85, June 1994. ISSN 0960-3174 (Print) 1573-1375 (Online). doi:10.1007/BF00175354. Online available at http://samizdat.mines.edu/ga_tutorial/ga_tutorial.ps and <http://www.citeulike.org/user/Bc91/article/1449453> [accessed 2007-08-12]. Also published as technical report [1920].
- [646] Chang Wook Ahn and R. S. Ramakrishna. Augmented compact genetic algorithm. In Roman Wyrzykowski, Jack Dongarra, Marcin Paprzycki, and Jerzy Wasniewski, editors, *PPAM 2003: Proceedings of 5th International Conference on Parallel Processing and Applied Mathematics, revised papers*, volume 3019(1–2)/2004 of *Lecture Notes in Computer Science (LNCS)*, pages 560–565. Springer Berlin / Heidelberg, September

- 2003, Czestochowa, Poland. ISBN 978-3-540-21946-0. doi:10.1007/b97218. Published 2004.
- [647] James E. Baker. Adaptive Selection Methods for Genetic Algorithms. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pages 101–111, 1985. In proceedings [751].
- [648] Tobias Blickle and Lothar Thiele. A Comparison of Selection Schemes Used in Genetic Algorithms. Technical Report 11, Computer Engineering and Communication Networks Lab (TIK), Swiss Federal Institute of Technology (ETH), Gloriastrasse 35, 8092 Zurich, Switzerland, 1995. Online available at <http://citeseer.ist.psu.edu/16412.html> and <http://www.tik.ee.ethz.ch/~tec/publications/bt95a/> [accessed 2007-07-28].
- [649] J. David Schaffer. *Multiple Objective Optimization with Vector Evaluated Genetic Algorithms*. PhD thesis, Vanderbilt University, 1984.
- [650] J. David Schaffer. Multiple Objective Optimization with Vector Evaluated Genetic Algorithms. In *Proceedings of the 1st International Conference on Genetic Algorithms and Their Applications*, pages 93–100, 1985. In proceedings [751].
- [651] Jon T. Richardson, Mark R. Palmer, Gunar E. Liepins, and Mike R. Hilliard. Some Guidelines for Genetic Algorithms with Penalty Functions. In *Proceedings of the third International Conference on Genetic Algorithms ICGA*, pages 191–197, 1989. In proceedings [652].
- [652] J. David Schaffer, editor. *Proceedings of the 3rd International Conference on Genetic Algorithms*, June 1989, George Mason University, Fairfax, Virginia, USA. Morgan Kaufmann, San Francisco, CA. ISBN 1-55860-066-3.
- [653] Carlos M. Fonseca and Peter J. Fleming. Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 416–423, 1993. In proceedings [748]. Online available at <http://citeseer.ist.psu.edu/fonseca93genetic.html> and <http://www.lania.mx/~ccoello/EM00/fonseca93.ps.gz> [accessed 2007-08-29].
- [654] Shinya Watanabe, Tomoyuki Hiroyasu, and Mitsunori Miki. NCGA: Neighborhood Cultivation Genetic Algorithm for Multi-Objective Optimization Problems. In *GECCO Late Breaking Papers; Late Breaking papers at the Genetic and Evolutionary Computation Conference (GECCO-2002)*, pages 458–465, 2002. In proceedings [564]. Online available at <http://citeseer.ist.psu.edu/595343.html> and <http://www.lania.mx/~ccoello/EM00/watanabe02.pdf.gz> [accessed 2007-07-29].
- [655] Nils Aaall Barricelli. Esempi numerici di processi di evoluzione. *Methodos*, pages 45–68, 1954.
- [656] Nils Aaall Barricelli. Symbiogenetic Evolution Processes Realized by Artificial Methods. *Methodos*, 9:35–36, 1957.
- [657] Nils Aaall Barricelli from the Department of Biology, Division of Molecular Biology, Vanderbilt University, Nashville, Tennessee. Numerical Testing of Evolution Theories. Part I. Theroetical Introduction and Basic Tests. *Acta Biotheoretica*, 16(1/2):69–98, March 1962. ISSN 0001-5342 (Print) 1572-8358 (Online). doi:10.1007/BF01556771. Received: 27 November 1961, See also [658]. Online available at <http://www.springerlink.com/content/y502315688024453/fulltext.pdf> [accessed 2007-10-31].
- [658] Nils Aaall Barricelli from the Department of Biology, Division of Molecular Biology, Vanderbilt University, Nashville, Tennessee. Numerical Testing of Evolution Theories. Part II. Preliminary Tests of Performance. Symbiogenesis and Terrestrial Life. *Acta Biotheoretica*, 16(3/4):99–126, September 1963. ISSN 0001-5342 (Print) 1572-8358 (Online). doi:10.1007/BF01556602. Received: 27 November 1961, See also [657]. Online available at <http://www.springerlink.com/content/h85817217u25w6q7/fulltext.pdf> [accessed 2007-10-31].
- [659] Alex S. Fraser. Simulation of Genetic Systems by Automatic Digital Computers. *Australian Journal of Biological Science*, 10:484–491, 1957. ISSN 0004-9417.

- [660] Hans J. Bremermann from the Department of Mathematics University of California Berkeley, California. Optimization Through Evolution and Recombination. *Self-Organizing systems*, pages 93–10, 1962. Online available at <http://holtz.org/Library/Natural%20Science/Physics/> [accessed 2007-10-31].
- [661] Woodrow “Woody” Wilson Bledsoe. Lethally Dependent Genes Using Instant Selection. Technical Report PRI 1, Panoramic Research Inc., Palo Alto, California, USA, 1961.
- [662] Woodrow “Woody” Wilson Bledsoe. The Use of Biological Concepts in the Analytical Study of Systems. Technical Report PRI 2, Panoramic Research, Inc., Palo Alto, California, USA, 1961. Presented at ORSA-TIMS National Meeting, San Francisco, California, November 10, 1961.
- [663] Woodrow “Woody” Wilson Bledsoe. An Analysis of Genetic Populations. Technical Report, Panoramic Research Inc., Palo Alto, California, USA, 1962.
- [664] Woodrow “Woody” Wilson Bledsoe. The Evolutionary Method in Hill Climbing: Convergence Rates. Technical Report, Panoramic Research Inc., Palo Alto, California, USA, 1962.
- [665] Woodrow “Woody” Wilson Bledsoe and I. Browning. Pattern Recognition and Reading by Machine. In *Proceedings of the Eastern Joint Computer Conference (EJCC)*, pages 225–232, 1959.
- [666] John Daniel Bagley. *The Behavior of Adaptive Systems which employ Genetic and Correlation Algorithms*. PhD thesis, The University of Michigan, College of Literature, Science, and the Arts, Computer and Communication Sciences Department, Ann Arbor, MI, USA, December 1967. Order No. AAI6807556. Published as Technical Report and in Dissertations International 28(12), 5106B, University Microfilms No. 68-7556. Online available at <http://hdl.handle.net/2027.42/3354> [accessed 2007-10-31].
- [667] Daniel Joseph Cavicchio, Jr. *Adaptive Search using Simulated Evolution*. PhD thesis, The University of Michigan, College of Literature, Science, and the Arts, Computer and Communication Sciences Department, Ann Arbor, Michigan, USA, August 1970. Published as Technical Report. Chairman: John Henry Holland. Online available at <http://hdl.handle.net/2027.42/4042> [accessed 2007-10-31].
- [668] Daniel Raymond Frantz. *Nonlinearities in Genetic Adaptive Search*. PhD thesis, The University of Michigan, Ann Arbor, MI, USA, September 1972. Order No. AAI7311116. Chairman: John Henry Holland. Published as Technical Report Nr. 138 and in Dissertations International 33(11), 5240B-5241B, University Microfilms No. 73-11116. Online available at <http://hdl.handle.net/2027.42/4950> [accessed 2007-10-31].
- [669] John Henry Holland. Outline for a Logical Theory of Adaptive Systems. *Journal of the ACM*, 9(3):297–314, 1962. ISSN 0004-5411. doi:10.1145/321127.321128. Online available at <http://portal.acm.org/citation.cfm?id=321128> [accessed 2007-07-28].
- [670] Jack L. Crosby. *Computer Simulation in Genetics*. John Wiley and Sons Ltd, January 1973. ISBN 0-4711-8880-8.
- [671] Carlos Artemio Coello Coello. An updated survey of GA-based multiobjective optimization techniques. *ACM Computing Surveys*, 32(2):109–143, 2000. ISSN 0360-0300. doi:10.1145/358923.358929. Online available at <http://citeseer.ist.psu.edu/coello98updated.html> and <http://portal.acm.org/citation.cfm?id=358923.358929> [accessed 2007-07-28].
- [672] Karl Sims. Artificial evolution for computer graphics. In *SIGGRAPH’91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, pages 319–328. ACM, New York, NY, USA, 1991. ISBN 0-89791-436-8. doi:10.1145/122718.122752. Online available at <http://doi.acm.org/10.1145/122718.122752> [accessed 2008-03-22].

- [673] Hideyuki Takagi from the Kyushi Institute of Design. Active User Intervention in an EC Search. In *Proceedings of International Conference on Information Sciences (JCIS2000)*, pages 995–998, 2000. In proceedings [547]. Online available at http://www.kyushu-id.ac.jp/~takagi/TAKAGI/IECpaper/JCIS2K_2.pdf [accessed 2007-08-29].
- [674] Hideyuki Takagi from the Kyushi Institute of Design. Interactive Evolutionary Computation: Fusion of the Capacities of EC Optimization and Human Evaluation. *Proceedings of the IEEE*, 89(9):1275–1296, 2001. Online available at <http://www.design.kyushu-u.ac.jp/~takagi/TAKAGI/IECsurvey.html> [accessed 2007-08-29].
- [675] Alex Kosorukoff. Human-based Genetic Algorithm. In *Proceedings of the 2001 IEEE International Conference on Systems, Man, and Cybernetics*, volume 5, pages 3464–3469, January 2001, Tucson, AZ, USA. ISBN 0-7803-7087-2. Also: IlliGAL Report No. 2001004, Online available at <http://citeseer.ist.psu.edu/kosorukoff01human.html> and <http://citeseer.ist.psu.edu/441929.html> [accessed 2007-07-28].
- [676] Michelle Okaley Hammond and Terence Claus Fogarty. Co-operative OuLiP-ian Generative Literature using Human Based Evolutionary Computing. In *Late breaking paper at Genetic and Evolutionary Computation Conference (GECCO'2005)*, 2005. In proceedings [556]. Distributed on CD-ROM at GECCO-2005. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2005lbp/papers/56-hammond.pdf> [accessed 2007-08-29].
- [677] David Levine. Application of a hybrid genetic algorithm to airline crew scheduling. *Computers & Operations Research*, 23(6):547–558, 1996. ISSN 0305-0548. Online available at <http://www.citeulike.org/user/ilapla/article/1443054> and <http://citeseer.ist.psu.edu/178394.html> [accessed 2007-07-29].
- [678] Gary A. Cleveland and Stephen F. Smith. Using Genetic Algorithms to Schedule Flow Shop Releases. In *ICGA, Proceedings of the third International Conference on Genetic Algorithms*, pages 160–169, 1989. In proceedings [652].
- [679] Raymond S. K. Kwan, Ann S. K. Kwan, and Anthony Wren. Driver scheduling using genetic algorithms with embedded combinatorial traits. In *Proceedings of 7th International Conference on Computer-Aided Scheduling of Public Transport*, volume 471, pages 81–102. Springer, Berlin, August 1997, Cambridge/Boston, MA, USA. ISBN 3-540-65775-4. Online available at <http://www.citeulike.org/user/ilapla/article/1443013> and <http://citeseer.ist.psu.edu/kwan97driver.html> [accessed 2007-07-29].
- [680] Steven J. Beaty. Genetic Algorithms for Instruction Sequencing and Scheduling. In *Workshop on Computer Architecture Technology and Formalism for Computer Science Research and Applications*, April 1992. Online available at <http://citeseer.ist.psu.edu/16699.html> and <http://emess.mscd.edu/~beaty/Dossier/Papers/italy.pdf> [accessed 2007-07-29].
- [681] Alain Cardon, Theirry Galinho, and Jean-Philippe Vacher. A Multi-Objective Genetic Algorithm in Job Shop Scheduling Problem to Refine an Agents' Architecture. In *Proceedings of EUROGEN'99*, 1999. In proceedings [526]. Online available at <http://www.jeo.org/emo/cardon99.ps.gz> [accessed <http://citeseer.ist.psu.edu/244765.html>]2008-04-05.
- [682] Alain Cardon, Theirry Galinho, and Jean-Philippe Vacher. An Agent Based Architecture for Job-Shop Scheduling Problem Using the Spirit of Genetic Algorithm. In *Proceedings of EUROGEN'99*, pages 12–19, 1999. In proceedings [526]. Online available at <http://www.mit.jyu.fi/eurogen99/papers/vacher12.ps> [accessed 2008-06-07].
- [683] Alain Cardon, Theirry Galinho, and Jean-Philippe Vacher. Using Genetic Algorithm in Job-Shop Scheduling Problem to Constraints Negotiators' Agents. In *Proceedings of EUROGEN'99*, pages 20–27, 1999. In proceedings [526].
- [684] António Gaspar Lopes da Cunha and José António Colaço Gomes Covas. RPSGAe – Reduced Pareto Set Genetic Algorithm: Application to Polymer Extrusion. In

- Metaheuristics for Multiobjective Optimisation*, pages 221–249. Springer, 2004. In collection [399].
- [685] Yong L. Xiao and Donald E. Williams. Game: Genetic Algorithm for Minimization of Energy, an Interactive Program for Three-dimensional Intermolecular Interactions. *Computers & Chemistry*, 18(2):199–201, 1994.
- [686] António Gaspar Lopes da Cunha. A Multi-Objective Evolutionary Algorithm for Solving Traveling Salesman Problems: Application to the Design of Polymer Extruders. In *Proceedings of the 7th International Conference on Adaptive and Natural Computing Algorithms, ICANNGA 2005, Part II*, pages 189–193, 2005. In proceedings [576].
- [687] António Gaspar Lopes da Cunha, José António Colaço Gomes Covas, and Pedro Oliveira. Optimization of Polymer Extrusion with Genetic Algorithms. *IMA Journal of Mathematics Applied in Business & Industry*, (9):267–277, 1998. Online available at <http://imaman.oxfordjournals.org/cgi/reprint/9/3/267.pdf> [accessed 2007-07-28].
- [688] D. M. Deaven and Kai-Ming Ho from the Physics Department, Ames Laboratory U.S. DOE, Ames, Iowa 50011, USA. Molecular geometry optimization with a genetic algorithm. *Physical Review Letters*, 75:288–291, July 1995. doi:10.1103/PhysRevLett.75.288. eprint arXiv:mr1-th/9506004. Online available at <http://adsabs.harvard.edu/abs/1995mr1.th...6004D> and http://prola.aps.org/abstract/PRL/v75/i2/p288_1 [accessed 2007-09-05].
- [689] Maria Wächtler. Entwicklung eines Programmes zur Auswertung von ⁵¹V-Festkörper-NMR-MAS-Spektren mit Hilfe genetischer Algorithmen am Beispiel von Spektren von Modellkomplexen für vanadiumhaltige Haloperoxidasen. Master's thesis, Friedrich-Schiller-Universität Jena, Chemisch-geowissenschaftliche Fakultät, Institut für Physikalische Chemie, 2007. Supervisors: Prof. Dr. Gerd Buntkowsky and Prof. Dr. Jürgen Popp.
- [690] R. S. Judson, M. E. Colvin, Juan C. Meza, A. Huffer, and D. Gutierrez from the Sandia National Laboratories, cent. computational eng., Livermore CA 94551-0969, USA. Do intelligent configuration search techniques outperform random search for large molecules? *International Journal of Quantum Chemistry*, 44(2):277–290, 1992. ISSN 0020-7608. Coden: IJQCB2.
- [691] Juan C. Meza and M. L. Martinez from the Scientific Computing Department, Sandia National Laboratories, Livermore, CA 94551-0969. On the Use of Direct Search Methods for the Molecular Conformation Problem. *Journal of Computational Chemistry*, 15(6):627–632, 1994. Online available at <http://crd.lbl.gov/~meza/papers/jcc.pdf> [accessed 2008-06-14].
- [692] A. Cagnoni, A. Dobrzeniecki, R. Poli, and J. Yanch. Genetic algorithm-based interactive segmentation of 3D medical images. *Image and Vision Computing*, 17(12):881–895, October 1999. Online available at <http://citeseer.ist.psu.edu/cagnoni99genetic.html> [accessed 2007-07-29].
- [693] Rafal Smigrodzki, Ben Goertzel, Cassio Pennachin, Lucio Coelho, Francisco Prodocimi, and W. Davis Parker Jr. Genetic algorithm for analysis of mutations in Parkinson's disease. *Artificial Intelligence in Medicine*, 35:227–241, November 2005. Online available at <http://dx.doi.org/10.1016/j.artmed.2004.11.006> [accessed 2007-08-05].
- [694] Hongmei Yan, Yingtao Jiang, Jun Zheng, Chenglin Peng, and Shouzhong Xiao. Discovering Critical Diagnostic Features for Heart Diseases with a Hybrid Genetic Algorithm. In Faramarz Valafar and Homayoun Valafar, editors, *Proceedings of the International Conference on Mathematics and Engineering Techniques in Medicine and Biological Sciences, METMBS'03*, pages 406–409. CSREA Press, June 2003, Las Vegas, Nevada, USA. ISBN 1-932415-04-1.

- [695] Staal A. Vinterbo and Lucila Ohno-Machado. A genetic algorithm approach to multi-disorder diagnosis. *Artificial Intelligence in Medicine*, 18(2):117–132, 2000. Online available at [http://dx.doi.org/10.1016/S0933-3657\(99\)00036-6](http://dx.doi.org/10.1016/S0933-3657(99)00036-6) [accessed 2007-09-05].
- [696] Hokey Min, Tomasz G. Smolinski, and Grzegorz M. Boratyn. A Genetic Algorithm-based Data Mining Approach to Profiling the Adopters And Non-Adopters of E-Purchasing. In W. W. Smari, editor, *Information Reuse and Integration, Third International Conference, IRI-2001*, pages 1–6. International Society for Computers and Their Applications (ISCA), 2001. Online available at <http://citeseer.ist.psu.edu/min01genetic.html> [accessed 2007-07-29].
- [697] Ali Kamrani, Wang Rong, and Ricardo Gonzalez. A genetic algorithm methodology for data mining and intelligent knowledge acquisition. *Computers & Industrial Engineering*, 40(4):361–377, September 2001. ISSN 0360-8352.
- [698] Janaki Gopalan, Reda Alhajj, and Ken Barker. Discovering Accurate and Interesting Classification Rules Using Genetic Algorithm. In Sven F. Crone, Stefan Lessmann, and Robert Stahlbock, editors, *Proceedings of the 2006 International Conference on Data Mining, DMIN 2006*, pages 389–395. CSREA Press, June 2006, Las Vegas, Nevada, USA. ISBN 1-60132-004-3. Online available at <http://ww1.ucmss.com/books/LFS/CSREA2006/DMI5509.pdf> [accessed 2007-07-29].
- [699] George G. Szpiro from the Israeli Centre for Academic Studies, Kiriath Ono, Israel. A Search for Hidden Relationships: Data Mining with Genetic Algorithms. *Computational Economics*, 10(3):267–277, August 1997. Online available at <http://citeseer.ist.psu.edu/459445.html> and <http://ideas.repec.org/a/kap/compec/v10y1997i3p267-77.html> [accessed 2007-07-03].
- [700] Jinxiang Chai and Songde Ma. Robust epipolar geometry estimation using genetic algorithm. *Pattern Recognition Letters*, 19(9):829–838, July 1998. ISSN 0167-8655. doi:10.1016/S0167-8655(98)00032-4. See also [701]. Online available at [http://dx.doi.org/10.1016/S0167-8655\(98\)00032-4](http://dx.doi.org/10.1016/S0167-8655(98)00032-4) [accessed 2008-03-22].
- [701] Jinxiang Chai and Songde Ma. Robust Epipolar Geometry Estimation Using Genetic Algorithm. In Roland T. Chin and Ting-Chuen Pong, editors, *ACCV, Proceedings of Computer Vision - ACCV'98, Third Asian Conference on Computer Vision, Volume I*, volume 1351 of *Lecture Notes in Computer Science (LNCS)*, pages 272–279. Springer, January 1998, Hong Kong, China. ISBN 3-540-63930-6. doi:10.1016/S0167-8655(98)00032-4. See also [700].
- [702] Ian Hsieh, Kiat-Choong Chen, and Cao An Wang. A Genetic Algorithm for the Minimum Tetrahedralization of a Convex Polyhedron. In *CCCG Proceedings of the 15th Canadian Conference on Computational Geometry, CCCG'03*, pages 115–119, August 2003, Halifax, Canada.
- [703] Sourav Kundu, Kazuto Seto, and Shigeru Sugino. Genetic Algorithm Based Design of Passive Elements for Vibration Control. In *Proceedings of 4th MOVIC Conference (Conference On Motion and Vibration Control)*, volume 3, pages 1183–1188, August 1998. Online available at <http://citeseer.ist.psu.edu/349539.html> and <http://en.scientificcommons.org/362076> [accessed 2007-08-13].
- [704] Sourav Kundu, Kazuto Seto, and Shigeru Sugino. Genetic Algorithm Application to Vibration Control of Tall Flexible Structures. In *Proceedings of The First IEEE International Workshop on Electronic Design, Test and Applications (DELTA'02)*, pages 333–337, January 29–31, 2002, Christchurch, New Zealand. IEEE Computer Society, Los Alamitos, CA, USA. ISBN 0-7695-1453-7. doi:10.1109/DELTA.2002.994641.
- [705] D. Yuret and M. Maza. A genetic algorithm system for predicting the OEX. *Technical Analysis of Stocks & Commodities*, pages 58–64, June 1994. Online available at <http://www.denizyuret.com/pub/tasc94.ps.gz> and <http://citeseer.ist.psu.edu/yuret94genetic.html> [accessed 2007-08-24].
- [706] Khaled El-Fakihi, Hirozumi Yamaguchiz, and Gregor von Bochmann. A Method and a Genetic Algorithm for Deriving Protocols for Distributed Applications

- with Minimum Communication Cost. In *Proceedings of Eleventh IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS'99)*, November 3–6, 1999, Cambridge, Boston, MA, USA. Online available at <http://citeseer.ist.psu.edu/430349.html> and <http://www-higashi.ist.osaka-u.ac.jp/~h-yamagu/resource/pdcs99.pdf> [accessed 2007-09-14].
- [707] Lidia A. R. Yamamoto and Christian F. Tschudin from the Computer Science Department, University of Basel, Bernoullistrasse 16, CH-4056 Basel, Switzerland. Genetic Evolution of Protocol Implementations and Configurations. In *IFIP/IEEE International workshop on Self-Managed Systems and Services (SelfMan 2005)*, 2005, Nice, France. Online available at <http://cn.cs.unibas.ch/pub/doc/2005-selfman.pdf> [accessed 2007-09-17].
- [708] Sancho Salcedo-Sanz and Xin Yao. Assignment of cells to switches in a cellular mobile network using a hybrid Hopfield network-genetic algorithm approach. *Applied Soft Computing*, 8(1):216–224, January 2008. ISSN 1568-4946. doi:10.1016/j.asoc.2007.01.002.
- [709] Sami Khuri and Teresa Chiu from the Department of Mathematics and Computer Science, San José State University, One Washington Square, San José, CA. Heuristic algorithms for the terminal assignment problem. In *SAC'97: Proceedings of the 1997 ACM symposium on Applied computing*, pages 247–251, 1997, San Jose, California, United States. ACM, New York, NY, USA. ISBN 0-89791-850-9. doi:10.1145/331697.331748. Online available at <http://doi.acm.org/10.1145/331697.331748> and <http://citeseer.ist.psu.edu/khuri97heuristic.html> [accessed 2008-06-15].
- [710] Nicholas Peter Sharples and Ian Wakeman from the University of Sussex Falmer, Brighton, BN1 9QH. Protocol Construction Using Genetic Search Techniques. In *Real-World Applications of Evolutionary Computing, EvoWorkshops 2000: EvoIASP, EvoSCONDI, EvoTel, EvoSTIM, EvoROB, and EvoFlight*, pages 73–95, 2000. doi:10.1007/3-540-45561-2.23. In proceedings [541].
- [711] Nicholas Peter Sharples from the University of Sussex Falmer, Brighton, BN1 9QH. *Evolutionary Approaches to Adaptive Protocol Design*. PhD thesis, School of Cognitive & Computing Sciences of the University of Sussex, August 2001. ASIN: B001ABO1DK.
- [712] Mark C. Sinclair from the Department of Electronic Systems Engineering, University of Essex, Wivenhoe Park, Colchester, C04 3SQ, UK. Minimum cost topology optimisation of the COST 239 European optical network. In *Proceedings of the International Conference on Artificial Neural Networks and Genetic Algorithms*, pages 26–29, 1995. In proceedings [581]. Online available at <http://citeseer.ist.psu.edu/63056.html> [accessed 2008-07-29].
- [713] L.G. Tan and Mark C. Sinclair from the Department of Electronic Systems Engineering, University of Essex, Wivenhoe Park, Colchester, C04 3SQ, UK. Wavelength Assignment between the Central Nodes of the COST 239 European Optical Network. In *Proceedings of the 11th UK Performance Engineering Workshop*, pages 235–247, September 1995, Liverpool, UK. Online available at <http://citeseer.ist.psu.edu/99015.html> and http://uk.geocities.com/markcsinclair/ps/ukpew11_tan.ps.gz [accessed 2008-07-29].
- [714] Anup Kumar, Rakesh M. Pathak, and M. C. Gupta. Genetic Algorithm based Approach for Designing Computer Network Topologies. In *CSC'93: Proceedings of the 1993 ACM conference on Computer science*, pages 358–365, 1993, Indianapolis, Indiana, United States. ACM, New York, NY, USA. ISBN 0-89791-558-5. doi:<http://doi.acm.org/10.1145/170791.170871>. Online available at <http://doi.acm.org/10.1145/170791.170871> [accessed 2008-08-01].
- [715] Samuel Pierre and Gisèle Legault from the LICEF Laboratory, Tél'e-université, Université du Québec, 1001, rue Sherbrooke Est. 4 étage, Case Postale 5250, Succ. C,

- Montréal, Québec, Canada, H2L 4L5. An evolutionary approach for configuring economical packet switched computer networks. *AI in Engineering*, 10(2):127–134, 1996. doi:10.1016/0954-1810(95)00022-4. Online available at [http://dx.doi.org/10.1016/0954-1810\(95\)00022-4](http://dx.doi.org/10.1016/0954-1810(95)00022-4) [accessed 2008-08-01].
- [716] Samuel Pierre and Gisèle Legault. A genetic algorithm for designing distributed computer network topologies. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 28(2):249–258, April 1998. doi:10.1109/3477.662766.
- [717] Alejandro Quintero and Samuel Pierre from the Department of Computer Engineering, Mobile Computing, and Networking Research (LARIM), École Polytechnique de Montréal, C.P. 6079, Succ. Centre-Ville, Montréal, Québec, Canada H3C 3A7. Assigning cells to switches in cellular mobile networks: a comparative study. *Computer Communications*, 26(9):950–960, June 2003. doi:10.1016/S0140-3664(02)00224-4. Online available at [http://dx.doi.org/10.1016/S0140-3664\(02\)00224-4](http://dx.doi.org/10.1016/S0140-3664(02)00224-4) [accessed 2008-08-01].
- [718] David Brittain, Jon Sims Williams, and Chris McMahon from the Department of Engineering Mathematics, University of Bristol, Bristol, UK. A Genetic Algorithm Approach To Planning The Telecommunications Access Network. In *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA97)*, pages 623–628, 1997. In proceedings [746]. Online available at <http://citeseer.ist.psu.edu/brittain97genetic.html> [accessed 2008-08-01].
- [719] King-Tim Ko, Kit-Sang Tang, Cheung-Yau Chan, Kim-Fung Man, and Sam Kwong. Using Genetic Algorithms to Design Mesh Networks. *Computer*, 30(8):56–61, August 1997. ISSN 0018-9162. doi:10.1109/2.607086.
- [720] Berna Dengiz, Fulya Altıparmak, and Alice E. Smith. Local Search Genetic Algorithm for Optimal Design of Reliable Networks. *IEEE Trans. on Evolutionary Computation*, 1(3):179–188, 1997. Online available at <http://citeseer.ist.psu.edu/dengiz97local.html> [accessed 2008-08-01].
- [721] F. Xiang, L. Junzhou, W. Jieyi, and G. Guanqun from the Department of Computer Science and Engineering, Southeast University, Nanjing 210096, China. QoS routing based on genetic algorithms. *Computer Communications*, 22(15-16):1392–1399, September 25, 1999. doi:10.1016/S0140-3664(99)00113-9. Online available at <http://neo.lcc.uma.es/opticomm/main.html> and [http://dx.doi.org/10.1016/S0140-3664\(99\)00113-9](http://dx.doi.org/10.1016/S0140-3664(99)00113-9) [accessed 2008-08-01].
- [722] Zhao-Xia Wang, Zeng-Qiang Chen, and Zhu-Zhi Yuan from the Department of Automation, Nankai University, Tianjin 300071, P.R. China. QoS routing optimization strategy using genetic algorithm in optical fiber communication networks. *Journal of Computer Science and Technology*, 19(2):213–217, March 2004. ISSN 1000-9000. CODEN: JCTEEM.
- [723] Liming Zhu, Roger L. Wainwright, and Dale A. Schoenefeld from the Mathematical and Computer Sciences Department, The University of Tulsa, 600 South College Avenue, Tulsa, OK 74104-3189, USA. A genetic algorithm for the point to multipoint routing problem with varying number of requests. In *Proceedings of the IEEE World Congress on Computational Intelligence, The 1998 IEEE International Conference on Evolutionary Computation*, pages 171–176, 1998. doi:10.1109/ICEC.1998.699496. In proceedings [503]. Online available at <http://euler.mcs.utulsa.edu/~rogerw/papers/Zhu-ICEC98.pdf> and <http://citeseer.ist.psu.edu/382506.html> [accessed 2008-07-21].
- [724] Pablo Galiasso and Roger L. Wainwright from the Mathematical and Computer Sciences Department, The University of Tulsa, 600 South College Avenue, Tulsa, OK 74104-3189, USA. A hybrid genetic algorithm for the point to multipoint routing problem with single split paths. In *SAC '01: Proceedings of the 2001 ACM symposium on Applied computing*, pages 327–332, 2001, Las Vegas, Nevada, United States. ACM, New York, NY, USA. ISBN 1-58113-287-5. doi:10.1145/372202.372354. Online

- available at <http://doi.acm.org/10.1145/372202.372354> and <http://citeseer.ist.psu.edu/451433.html> [accessed 2008-07-21].
- [725] Jason D. Lohn and Silvano P. Colombano. A Circuit Representation Technique for Automated Circuit Design. *IEEE Transactions on Evolutionary Computation (IEEE-EC)*, 3(3):205, September 1999. Online available at <http://ic.arc.nasa.gov/people/jlohn/bio.html> and <http://citeseer.ist.psu.edu/lohn99circuit.html> [accessed 2007-08-07].
 - [726] Jason D. Lohn, Silvano P. Colombano, Garyl L. Haith, and Dimitris Stassinopoulos. A Parallel Genetic Algorithm for Automated Electronic Circuit Design. In *Proceedings of the Computational Aerosciences Workshop*, February 2000, NASA Ames Research Center. Online available at <http://ic.arc.nasa.gov/people/jlohn/bio.html> and <http://citeseer.ist.psu.edu/lohn00parallel.html> [accessed 2007-08-07].
 - [727] Jason D. Lohn, Garyl L. Haith, Silvano P. Colombano, and Dimitris Stassinopoulos. Towards Evolving Electronic Circuits for Autonomous Space Applications. In *Proceedings of the 2000 IEEE Aerospace Conference*, March 2000, Big Sky, MT. Online available at <http://ic.arc.nasa.gov/people/jlohn/bio.html> and <http://citeseer.ist.psu.edu/336276.html> [accessed 2007-08-07].
 - [728] Jarmo T. Alander. An Indexed Bibliography of Genetic Algorithms in Signal and Image Processing. Technical Report 94-1-SIGNAL, Department of Information Technology and Production Economics, University of Vaasa, P.O. Box 700, FIN-65101, Finland, May 18, 1998. Online available at <http://citeseer.ist.psu.edu/319835.html> and <ftp://ftp.uwasa.fi/cs/report94-1/gaSIGNAlbib.ps.Z> [accessed 2008-05-18].
 - [729] Heinz Mühlenbein from the GMD Schloss Birlinghoven, D-5205 Sankt Augustin 1, Germany. Parallel Genetic Algorithms in Combinatorial Optimization. In Osman Balci, Ramesh Sharda, and Stavros A. Zenios, editors, *Selected papers from the Conference Computer Science and Operations Research: New Developments in Their Interfaces*, pages 441–456, January 8–10, 1992, Williamsburg, Virginia. Pergamon Press, Oxford. ISBN 0080408060, 978-0080408064. Online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.78.990> [accessed 2008-06-23].
 - [730] Sami Khuri, Thomas Bäck, and Jörg Heitkötter. The zero/one multiple knapsack problem and genetic algorithms. In *Proceedings of the 1994 ACM Symposium on Applied Computing, SAC'94*, pages 188–193, March 6–8, 1994, Phoenix, Arizona, USA. ACM, New York, NY, USA. ISBN 0-89791-647-6. doi:10.1145/326619.326694. Online available at <http://doi.acm.org/10.1145/326619.326694> and <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.40.4393> [accessed 2008-06-23].
 - [731] Christopher R. Stephens, Marc Toussaint, Darrell L. Whitley, and Peter F. Stadler, editors. *Revised Selected Papers of the 9th International Workshop on Foundations of Genetic Algorithms IX, FOGA 2007*, volume 4436/2007 of *Lecture Notes in Computer Science (LNCS)*, January 8–11, 2007, Ciudad Universitaria ("University City"), Universidad Nacional Autonoma de Mexico, Mexico City, México. Springer, Berlin Heidelberg. ISBN 978-3-540-73479-6. doi:10.1007/978-3-540-73482-6. see <http://www.sigevo.org/foga-2007/index.html> [accessed 2007-09-01].
 - [732] Alden H. Wright, Michael D. Vose, Kenneth Alan De Jong, and Lothar M. Schmitt, editors. *Revised Selected Papers of the 8th International Workshop on Foundations of Genetic Algorithms, FOGA 2005*, volume 3469/2005 of *Lecture Notes in Computer Science (LNCS)*, January 5–9, 2005, Aizu-Wakamatsu City, Japan. Springer Verlag, Berlin Heidelberg. ISBN 978-3-540-27237-3. doi:10.1007/11513575_1. Published August 22, 2005. see <http://www.cs.unt.edu/foga05/> [accessed 2007-09-01].
 - [733] Kenneth Alan De Jong, Riccardo Poli, and Jonathan E. Rowe, editors. *Foundations of Genetic Algorithms 7*, September 4–6, 2002, Torremolinos, Spain. Morgan Kaufmann, San Mateo, CA, USA. ISBN 0-12-208155-2. Published 2003.
 - [734] William M. Spears and Worthy N. Martin, editors. *Proceedings of the Sixth Workshop on Foundations of Genetic Algorithms*, July 21–23, 2000, Charlottesville, VA, USA.

- Morgan Kaufmann, San Mateo, CA, USA. ISBN 1-55860-734-X. see <http://www.cs.uwo.edu/~wspears/foga00/index.html> [accessed 2007-09-01].
- [735] Wolfgang Banzhaf and Colin R. Reeves, editors. *Proceedings of the Fifth Workshop on Foundations of Genetic Algorithms*, July 22–25, 1998, Madison, WI, USA. Morgan Kaufmann, San Mateo, CA, USA. ISBN 1-55860-559-2. Published April 2, 1991.
 - [736] Richard K. Belew and Michael D. Vose, editors. *Proceedings of the 4th Workshop on Foundations of Genetic Algorithms*, August 5, 1996, University of San Diego, San Diego, CA, USA. Morgan Kaufmann, San Francisco, California, USA. ISBN 1-55860-460-X. Published March 1, 1997.
 - [737] L. Darrell Whitley and Michael D. Vose, editors. *Proceedings of the Third Workshop on Foundations of Genetic Algorithms*, July 31–August 2, 1994, Estes Park, Colorado, USA. Morgan Kaufmann, San Francisco, CA, USA. ISBN 1-55860-356-5. Published June 1, 1995.
 - [738] L. Darrell Whitley, editor. *Proceedings of the Second Workshop on Foundations of Genetic Algorithms*, July 26–29, 1992, Vail, Colorado, USA. Morgan Kaufmann, San Mateo, CA, USA. ISBN 1-55860-263-1. Published February 1, 1993.
 - [739] Jarmo Tapani Alander, editor. *Proceedings of the Third Nordic Workshop on Genetic Algorithms (3NWGA)*, August 20–22, 1997, Helsinki, Finland. Finnish Artificial Intelligence Society (FAIS). Online available at <ftp://ftp.uwasa.fi/cs/3NWGA> [accessed 2008-07-24].
 - [740] Jarmo Tapani Alander, editor. *Proceedings of the Second Nordic Workshop on Genetic Algorithms (2NWGA)*, Proceedings of the University of Vaasa, Nro. 13, August 19–23, 1996, Vaasa, Finland. Department of Information Technology and Production Economics, University of Vaasa, P.O. Box 700, FIN-65101 Vaasa, Finland. In collaboration with the Finnish Artificial Intelligence Society. Online available at <ftp://ftp.uwasa.fi/cs/2NWGA> [accessed 2008-07-24].
 - [741] Jarmo Tapani Alander, editor. *Proceedings of the First Nordic Workshop on Genetic Algorithms (1NWGA, The Third Finnish Workshop on Genetic Algorithms, 3FWGA)*, Proceedings of the University of Vaasa, Nro. 2, June 9–13, 1995, Vaasa, Finland. Department of Information Technology and Production Economics, University of Vaasa, P.O. Box 700, FIN-65101 Vaasa, Finland. Technical Report 95-1. Online available at <ftp://ftp.uwasa.fi/cs/1NWGA> [accessed 2008-07-24].
 - [742] Jarmo Tapani Alander, editor. *Proceedings of the Second Finnish Workshop on Genetic Algorithms and Their Applications (2FWGA)*, March 16–18, 1994, Vaasa, Finland. Department of Information Technology and Production Economics, University of Vaasa, P.O. Box 700, FIN-65101 Vaasa, Finland. ISBN 951-683-526-0. Technical Report 94-2. Online available at <ftp://ftp.uwasa.fi/cs/report94-2> [accessed 2008-07-24].
 - [743] Jarmo Tapani Alander, editor. *Proceedings of the 1st Finnish Workshop on Genetic Algorithms and Their Applications*, November 4–5, 1992, Espoo, Finland. Department of Computer Science, Helsinki University of Technology, Helsinki, Finland. Technical Report TKO-A30, partly in Finnish, published 1993.
 - [744] *Proceedings of the 2nd IEE/IEEE International GALEZIA Conference*, September 2–4, 1997, University of Strathclyde, Glasgow, UK. IEE. ISBN 0-85296-693-8.
 - [745] A. M. S. Zalzala, editor. *First International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications (GALESIA)*, volume 414, September 12–14, 1995, Scheffield, UK. IEE Conference Publication, Institution of Engineering and Technology. ISBN 978-0852966501.
 - [746] Thomas Bäck, editor. *Proceedings of The Seventh International Conference on Genetic Algorithms ICGA '97*, July 19–23, 1997, Michigan State University, East Lansing, Michigan, USA. Morgan Kaufmann Publishers, San Francisco, CA, USA. ISBN 1-55860-487-1. see <http://garage.cse.msu.edu/icga97/> [accessed 2007-09-01].

- [747] Larry J. Eshelman, editor. *Proceedings of the Sixth International Conference on Genetic Algorithms*, July 15–19, 1995, Pittsburgh, PA, USA. Morgan Kaufmann, San Francisco, CA. ISBN 1-55860-370-0.
- [748] Stephanie Forrest, editor. *Proceedings of the 5th International Conference on Genetic Algorithms*, June 1993, Urbana-Champaign, IL, USA. Morgan Kaufmann, San Francisco, CA. ISBN 1-55860-299-2.
- [749] Richard K. Belew and Lashon Bernard Booker, editors. *Proceedings of the 4th International Conference on Genetic Algorithms*, July 1991, San Diego, CA, USA. Morgan Kaufmann, San Francisco, CA. ISBN 1-55860-208-9.
- [750] John J. Grefenstette, editor. *Proceedings of the 2nd International Conference on Genetic Algorithms and their Applications*, July 1987, Cambridge, Massachusetts, USA. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, USA. ISBN 0-8058-0158-8.
- [751] John J. Grefenstette, editor. *Proceedings of the 1st International Conference on Genetic Algorithms and their Applications*, July 24–26, 1985, Carnegie-Mellon University, Pittsburgh, PA, USA. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, USA. ISBN 0-8058-0426-9.
- [752] Lawrence Davis, editor. *Handbook of Genetic Algorithms*. Thomson Publishing Group / Van Nostrand Reinhold Company, January 1991. ISBN 0442001738, 978-0442001735.
- [753] Randy L. Haupt and Sue Ellen Haupt. *Practical Genetic Algorithms*. Wiley-Interscience, December 19, 1997. ISBN 0471188735, 978-0471188735, 2nd ed: 978-0-471-45565-3. Second edition: June 2004, partly online available at http://books.google.de/books?id=W064i0g_abIC [accessed 2008-06-08].
- [754] Mitsuo Gen and Runwei Cheng. *Genetic Algorithms and Engineering Design*. Wiley-Interscience, January 7, 1997. ISBN 0471127418, 978-0471127413. Partly online available at <http://books.google.de/books?id=MCHCaJAHFJAC> [accessed 2008-04-01].
- [755] Lance D. Chambers, editor. *Practical Handbook of Genetic Algorithms: Applications*, volume I of *Practical Handbook of Genetic Algorithms*. Chapman & Hall/CRC Press, Inc., Boca Raton, FL, USA, 2nd ed: 2000 edition, 1995. ISBN 0849325196, 1584882409, 978-1584882404.
- [756] Lance D. Chambers, editor. *Practical Handbook of Genetic Algorithms: New Frontiers*, volume II of *Practical Handbook of Genetic Algorithms*. CRC Press, Inc., Boca Raton, FL, USA, 1995. ISBN 0849325293. Partly online available at <http://books.google.de/books?id=9RCE3pgj9K4C> [accessed 2008-07-19].
- [757] Lance D. Chambers, editor. *Practical Handbook of Genetic Algorithms: Complex Coding Systems*, volume III of *Practical Handbook of Genetic Algorithms*. CRC Press, Inc., Boca Raton, FL, USA, December 23, 1998. ISBN 0849325390, 978-0849325397.
- [758] Mitsuo Gen and Runwei Chen. *Genetic Algorithms (Engineering Design and Automation)*. Wiley Series in Engineering Design and Automation. John Wiley and Sons Ltd, February 2001. ISBN 978-0-471-31531-5.
- [759] Erick Cant'u-Paz. *Efficient and Accurate Parallel Genetic Algorithms*, volume 1 of *Genetic Algorithms and Evolutionary Computation*. Springer, December 15, 2000. ISBN 978-0-7923-7221-9, 0792372212.
- [760] Jochen Heistermann. *Genetische Algorithmen. Theorie und Praxis evolutionärer Optimierung*. Teubner Verlag, 1994. ISBN 3815420571, 978-3815420577.
- [761] Eberhard Schöneburg, Frank Heinzmann, and Sven Feddersen. *Genetische Algorithmen und Evolutionsstrategien*. Addison Wesley Verlag, 1993. ISBN 3893194932, 978-3893194933.
- [762] Tomasz Dominik Gwiazda. *Crossover for single-objective numerical optimization problems*, volume 1 of *Genetic Algorithms Reference*. Lomianki, e-book: Tomasz Dominik Gwiazda, May 2006. ISBN 83-923958-3-2.
- [763] Robert Schaefer and Henryk Telega. *Foundations of Global Genetic Optimization*, volume 74/2007 of *Studies in Computational Intelligence*. Springer Berlin / Heidel-

- berg, 2007. ISBN 978-3-540-73191-7. doi:10.1007/978-3-540-73192-4. Series editor: Janusz Kacprzyk.
- [764] Charles L. Karr and L. Michael Freeman, editors. *Industrial Applications of Genetic Algorithms*. International Series on Computational Intelligence. CRC Press, Boca Raton, FL, December 29, 1998. ISBN 0849398010, 978-0849398018.
- [765] Lawrence Davis, editor. *Genetic Algorithms and Simulated Annealing*. Research Notes in Artificial Intelligence. Morgan Kaufmann Publishers Inc., Los Altos/San Francisco, CA, USA, October 1987. ISBN 0934613443, 978-0934613446.
- [766] Sanza Kazadi. Conjugate Schema and Basis Representation of Crossover and Mutation. *Evolutionary Computation*, 6(2):129–160, September 1998. Online available at http://www.jisan.org/research_collaborators/former_research/conjugate_schemata/papers/CONJUGATE98.PDF [accessed 2007-08-29].
- [767] Richard A. Caruana and J. David Schaffer. Representation and hidden bias: Gray vs. binary coding for genetic algorithms. In John E. Laird, editor, *Machine Learning, Proceedings of 5th International Conference on Machine Learning*, pages 153–161, June 1988, Ann Arbor, Michigan, USA. Morgan Kaufmann, San Mateo, California. ISBN 0-934613-64-8.
- [768] James D. Watson, Tania A. Baker, Stephen P. Bell, Alexander Gann, Michael Levine, and Richard Losick. *Molecular Biology of the Gene*. Benjamin Cummings, fifth (december 3, 2003) edition, 1987. ISBN 978-0805346350.
- [769] J. Peter Gogarten and Elena Hilario. Inteins, introns, and homing endonucleases: recent revelations about the life cycle of parasitic genetic elements. *BMC Evolutionary Biology*, 6(94), 2006. doi:10.1186/1471-2148-6-94. Online available at <http://dx.doi.org/10.1186/1471-2148-6-94> and <http://www.biomedcentral.com/content/pdf/1471-2148-6-94.pdf> [accessed 2008-03-23].
- [770] Walter Gilbert. Why genes in pieces. *Nature*, 271(5645):501, February 9, 1978. doi:10.1038/271501a0.
- [771] Claudia Crosio, Francesco Cecconi, Paolo Mariottini, Gianni Cesareni, Sydney Brenner, and Francesco Amaldi. Fugu intron oversize reveals the presence of U15 snoRNA coding sequences in some introns of the ribosomal protein S3 gene. *Genome Research*, 6(12):1227–1231, December 1996. ISSN 1088-9051. Online available at <http://www.genome.org/cgi/content/abstract/6/12/1227> [accessed 2008-03-23].
- [772] Nicol N. Schraudolph and Richard K. Belew from the Computer Science & Engineering Department, University of California, San Diego, La Jolla, CA, 92093-0114. Dynamic Parameter Encoding for Genetic Algorithms. *Machine Learning*, 9:9–21, June 1992. ISSN 0885-6125 (Print) 1573-0565 (Online). doi:10.1023/A:1022624728869. Online available at <http://citeseer.ist.psu.edu/schraudolph92dynamic.html> and <http://www.springerlink.com/content/u684071468r0x423/fulltext.pdf> [accessed 2007-08-29].
- [773] Steven Manos, Leon Poladian, Peter J. Bentley, and Maryanne Large. A Genetic Algorithm with a Variable-Length Genotype and Embryogeny for Microstructured Optical Fibre Design. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1721–1728, 2006. In proceedings [553]. Online available at <http://portal.acm.org/citation.cfm?id=1144278> [accessed 2007-08-17].
- [774] Kenneth O. Stanley and Risto Miikkulainen. A Taxonomy for Artificial Embryogeny. *Artificial Life*, 9(2):93–130, 2003. ISSN 1064-5462. Online available at <http://nn.cs.utexas.edu/downloads/papers/stanley.alife03.pdf> [accessed 2007-08-17].
- [775] Gregory Beurier, Fabien Michel, and Jacques Ferber. A morphogenesis model for multiagent embryogeny. In *ALIFE X, Tenth International Conference on the Simulation and Synthesis of Living Systems*, June 2006. Bloomington, Indiana, USA. Online available at <http://leri.univ-reims.fr/~fmichel/publi/pdfs/beurier06alifeX.pdf> [accessed 2007-08-17].

- [776] Sanjeev Kumar and Peter J. Bentley. Computational embryology: past, present and future. In *Advances in evolutionary computing: theory and applications*, pages 461–477. Springer, 2003. In collection [616]. Online available at <http://www.cs.ucl.ac.uk/staff/P.Bentley/KUBECH1.pdf> [accessed 2007-08-17].
- [777] Chris P. Bowers. Simulating Evolution with a Computational Model of Embryogeny: Obtaining Robustness from Evolved Individuals. In *Advances in Artificial Life, Proceedings of 8th European Conference*, pages 149–158, 2005. In proceedings [1921]. Online available at http://www.cs.bham.ac.uk/~cpb/publications/ecal05_bowers.pdf [accessed 2007-08-17].
- [778] Gwoing Tina Yu and Peter Bentley. Methods to Evolve Legal Phenotypes. In *PPSN V: Proceedings of the 5th International Conference on Parallel Problem Solving from Nature*, pages 280–291, 1998. doi:10.1007/BFb0056843. In proceedings [588]. Online available at <http://www.cs.ucl.ac.uk/staff/p.bentley/YUBEC2.pdf> [accessed 2007-08-17].
- [779] Aesop Rock. The Tugboat Complex pt. 3. In *Labor Days*, Aesop Rock, Omega One, and Blockhead, producers. Definitive Jux, September 18, 2001.
- [780] John J. Grefenstette. Deception Considered Harmful. In *Proceedings of the Second Workshop on Foundations of Genetic Algorithms*, pages 75–91, 1992. In proceedings [738]. Online available at <http://citeseer.ist.psu.edu/grefenstette92deception.html> [accessed <http://www.citeulike.org/user/pduval/article/1433820>]2007-08-12.
- [781] Melanie Mitchell, Stephanie Forrest, and John Henry Holland. The Royal Road for Genetic Algorithms: Fitness Landscapes and GA Performance. In *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 245–254, 1991. In proceedings [1922]. Online available at <http://citeseer.ist.psu.edu/mitchell91royal.html> and <http://web.cecs.pdx.edu/~mm/ecal92.pdf> [accessed 2007-10-15].
- [782] Peter John Angeline from the Natural Selection, Inc. A Historical Perspective on the Evolution of Executable Structures. *Fundamenta Informaticae*, 35(1–4):179–195, August 1998. ISSN 0169-2968. Special volume: Evolutionary Computation. Online available at <http://citeseer.ist.psu.edu/angeline98historical.html> and <http://www.natural-selection.com/Library/1998/gphist.pdf> [accessed 2008-05-29]. See also [782].
- [783] Richard M. Friedberg. A learning machine: Part I. *IBM Journal of Research and Development*, 2:2–13, November 1958. Online available at <http://www.research.ibm.com/journal/rd/021/ibmrd0201B.pdf> [accessed 2007-09-06]. See also [784].
- [784] Richard M. Friedberg, B. Dunham, and J. H. North. A Learning Machine: Part II. *IBM Journal of Research and Development*, 3(3):282–287, March 1959. Online available at <http://www.research.ibm.com/journal/rd/033/ibmrd0303H.pdf> [accessed 2007-09-06]. See also [783].
- [785] Lawrence J. Fogel, Alvin J. Owens, and Michael J. Walsh. *Artificial Intelligence through Simulated Evolution*. John Wiley & Sons, New York, USA, October 1966. ISBN 978-0471265160.
- [786] Stephen Frederick Smith. *A Learning System based on Genetic Adaptive Algorithms*. PhD thesis, University of Pittsburgh, Pittsburgh, PA, USA, 1980. Order No. AAI8112638, University Microfilms No. 81-12638.
- [787] Richard Forsyth from the Department of Mathematics, Polytechnic of North London, Holloway Road, London N78DB, UK. BEAGLE A Darwinian Approach to Pattern Recognition. *Kybernetes*, 10:159–166, 1981. Received December 17, 1980. Online available at http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/kybernetes_forsyth.pdf [accessed 2007-11-01] (copy from British Library May 1994).

- [788] Richard Forsyth and Roy Rada. *Machine Learning applications in Expert Systems and Information Retrieval*. Ellis Horwood series in Artificial Intelligence. Ellis Horwood, Chichester, UK, 1986. ISBN 0-7458-0045-9, 0-470-20309-9. Contains chapters on BEAGLE, see [787]. Also published by John Wiley & Sons Australia (May 28, 1986) and Halsted Press, New York, NY, USA.
- [789] Richard Forsyth. The Evolution of Intelligence. In Richard Forsyth, editor, *Machine Learning, Principles and Techniques*, chapter 4, pages 65–82. Chapman and Hall, 1989. ISBN 0-412-30570-4, 0-412-30580-1. Refers also to PC/BEAGLE, see [787].
- [790] Michael Lynn Cramer. A representation for the Adaptive Generation of Simple Sequential Programs. In *Proceedings of the 1st International Conference on Genetic Algorithms and their Applications*, pages 183–187, 1985. In proceedings [751]. Online available at <http://www.sover.net/~michael/nlc-publications/icga85/index.html> [accessed 2007-09-06].
- [791] Jürgen Schmidhuber. Evolutionary principles in self-referential learning. (On learning how to learn: The meta-meta-... hook.). Master’s thesis, Institut für Informatik, Technische Universität München, Munich, Germany, May 14, 1987. Online available at <http://www.idsia.ch/~juergen/diploma.html> [accessed 2007-11-01].
- [792] Dirk Dickmanns, Jürgen Schmidhuber, and Andreas Winklhofer. Der genetische Algorithmus: Eine Implementierung in Prolog. Fortgeschrittenenpraktikum, Institut für Informatik, Lehrstuhl Professor Radig, Technische Universität München, Munich, Germany, 1987. Online available at <http://www.idsia.ch/~juergen/geneticprogramming.html> [accessed 2007-11-01].
- [793] Joseph F. Hicklin. Application of the Genetic Algorithm to Automatic Program Generation. Master’s thesis, Department of Computer Science, University of Idaho, Moscow, ID, USA, 1986.
- [794] Cory Fujuki. An Evaluation of Hollands Genetic Algorithm Applied to a Program Generator. Master’s thesis, Department of Computer Science, University of Idaho, Moscow, ID, USA, 1986.
- [795] Cory Fujiko and John Dickinson. Using the genetic algorithm to generate LISP source code to solve the prisoner’s dilemma. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 236–240, 1987. In proceedings [750].
- [796] Arthur S. Bickel and Riva W. Bickel. Tree structured rules in genetic algorithms. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 77–81, 1987. In proceedings [750].
- [797] John R. Koza. *Non-Linear Genetic Algorithms for Solving Problems*. United States Patent and Trademark Office, 1988. United States Patent 4,935,877. Filed May 20, 1988. Issued June 19, 1990. Australian patent 611,350 issued september 21, 1991. Canadian patent 1,311,561 issued december 15, 1992.
- [798] John R. Koza. Hierarchical genetic algorithms operating on populations of computer programs. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence IJCAI-89*, pages 768–774, 1989. In proceedings [1923]. Partly online available at <http://dli.iiit.ac.in/ijcai/IJCAI-89-VOL1/PDF/123.pdf> [accessed 2008-05-29].
- [799] John R. Koza. The genetic programming paradigm: Genetically breeding populations of computer programs to solve problems. Technical Report STAN-CS-90-1314, Computer Science Department, Stanford University, Margaret Jacks Hall, Stanford, CA 94305, June 1990. Online available at <http://citeseer.ist.psu.edu/koza90genetic.html> [accessed 2008-05-29]. See also [1912].
- [800] John R. Koza from the Computer Science Department, Stanford University, Stanford CA 94305, USA. A Hierarchical Approach to Learning the Boolean Multiplexer Function. In *Proceedings of the First Workshop on Foundations of Genetic Algorithms*, pages 171–191, 1990. In proceedings [188]. Online available

- at <http://citeseer.ist.psu.edu/koza91hierarchical.html> and <http://books.google.de/books?id=Df12yLrlUZYC> [accessed 2008-05-29].
- [801] John R. Koza from the Computer Science Department, Margaret Jacks Hall, Stanford University, Stanford, CA 94305 USA. Genetic evolution and co-evolution of computer programs. In *Artificial Life II*, pages 603–629, 1990. Revised November 29, 1990. In proceedings [1924]. Online available at <http://citeseer.ist.psu.edu/177879.html> [accessed 2008-05-29].
 - [802] John R. Koza. Evolution and co-evolution of computer programs to control independent-acting agents. In *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 366–375, 1990. In proceedings [1925]. Online available at <http://www.genetic-programming.com/jkpdf/sab1990.pdf> [accessed 2008-05-29].
 - [803] John R. Koza. *Non-Linear Genetic Algorithms for Solving Problems by Finding a Fit Composition of Functions*. United States Patent and Trademark Office, August 1992. United States Patent 5,136,686. Filed March 28, 1990, Issued August 4, 1992.
 - [804] Roberto R. F. Mendes, Fabricio de B. Voznika, Alex A. Freitas, and Julio C. Nievola. Discovering fuzzy classification rules with genetic programming and co-evolution. In Luc de Raedt and Arno Siebes, editors, *Proceedings of 5th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'01)*, volume 2168 of *Lecture Notes in Computer Science (LNCS)*, subseries *Lecture Notes in Artificial Intelligence (LNAI)*, pages 314–325. Springer Verlag Berlin/Heidelberg, September 3–5, 2001, Freiburg, Germany. ISBN 978-3-540-42534-2. See also [805]. Online available at http://www.cs.kent.ac.uk/people/staff/aaf/pub_papers.dir/PKDD-2001.ps [accessed 2007-09-09].
 - [805] Roberto R. F. Mendes, Fabricio de B. Voznika, Alex A. Freitas, and Julio C. Nievola. Discovering fuzzy classification rules with genetic programming and co-evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, 2001. In proceedings [568] and [804]. Online available at <http://citeseer.ist.psu.edu/521000.html> [accessed 2007-09-09].
 - [806] John R. Koza. Concept formation and decision tree induction using the genetic programming paradigm. In *Parallel Problem Solving from Nature - Proceedings of 1st Workshop, PPSN 1*, pages 124–128, 1990. In proceedings [592]. Online available at <http://citeseer.ist.psu.edu/61578.html> [accessed 2007-09-09].
 - [807] Tatiana Kalganova and Julian Francis Miller. Evolving More Efficient Digital Circuits by Allowing Circuit Layout Evolution and Multi-Objective Fitness. In *Evolvable Hardware, Proceedings of 1st NASA/DoD Workshop on Evolvable Hardware (EH'99)*, page 54. IEEE Computer Society, July 1999, Pasadena, CA, USA. ISBN 0-7695-0256-3. Online available at <http://citeseer.ist.psu.edu/kalganova99evolving.html> [accessed 2007-09-09].
 - [808] Joel Jones from the Department of Computer Science, University of Alabama. Abstract Syntax Tree Implementation Idioms. In *Proceedings of The 10th Conference on Pattern Languages of Programs PLoP'2003*, September 8–12, 2003, Monticello, Illinois, USA. A workshop presentation online available at <http://hillside.net/plop/plop2003/Papers/Jones-ImplementingASTs.pdf> [accessed 2007-07-03].
 - [809] Nicola Howarth. Abstract Syntax Tree Design. Technical Report 23/08/95 APM.1551.01, Architecture Projects Management Limited, Poseidon House, Castle Park, Cambridge CB3 0RD, UK, August 23, 1995. Online available at <http://www.ansa.co.uk/ANSATech/95/Primary/155101.pdf> [accessed 2007-07-03].
 - [810] Robert Ian McKay, Xuan Hoai Nguyen, Peter Alexander Whigham, and Yin Shan. Grammars in Genetic Programming: A Brief Review. In L. Kang, Z. Cai, and Y. Yan, editors, *Progress in Intelligence Computation and Intelligence: Proceedings of the International Symposium on Intelligence, Computation and Applications*, pages 3–

18. China University of Geosciences Press, April 2005. Online available at <http://sc.snu.ac.kr/PAPERS/isica05.pdf> [accessed 2007-08-15].
- [811] Douglas A. Augusto and Helio J. C. Barbosa. Symbolic Regression via Genetic Programming. In *Proceedings of Sixth Brazilian Symposium on Neural Networks (SBRN'00)*, pages 173–178, 2000. IEEE Computer Society, Los Alamitos, CA, USA. ISBN 0-7695-0856-1.
- [812] Shengwu Xiong, Weiwu Wang, and Feng Li. A New Genetic Programming Approach in Symbolic Regression. In *Proceedings of 15th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'03)*, pages 161–167, November 3–5, 2003. IEEE Computer Society, Los Alamitos, CA, USA. doi:10.1109/TAI.2003.1250185.
- [813] Günther R. Raidl. A Hybrid GP Approach for Numerically Robust Symbolic Regression. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 323–328, 22–25 1998. In proceedings [913]. Online available at <http://citeseer.ist.psu.edu/549723.html> and <http://www.ads.tuwien.ac.at/publications/bib/pdf/raidl-98c.pdf> [accessed 2007-09-14].
- [814] Arturo Hernández Aguirre, Bill P. Buckles, and Carlos Artemio Coello Coello. A Genetic Programming Approach to Logic Function Synthesis by Means of Multiplexers. In *Evolvable Hardware, Proceedings of 1st NASA/DoD Workshop on Evolvable Hardware (EH'99)*, pages 46–53. IEEE Computer Society, July 1999, Pasadena, CA, USA. ISBN 0-7695-0256-3. Online available at <http://citeseer.ist.psu.edu/521808.html> [accessed 2007-09-09].
- [815] Francisco Fernandez de Vega. *Distributed Genetic Programming Models with Application to Logic Synthesis on FPGAs*. PhD thesis, University of Extremadura, 2001. Online available at <http://cum.unex.es/profes/profes/fcofdez/escritorio/investigacion/pgp/thesis/phd.html> [accessed 2007-09-09]. Spanish version: [1926].
- [816] Faizad Javed, Barrett R. Bryant, M. Črepinček, Marjan Mernik, and Alan Sprague. Context-free grammar induction using genetic programming. In *ACM-SE 42: Proceedings of the 42nd annual Southeast regional conference*, pages 404–405, 2004, Huntsville, Alabama. ACM Press, New York, NY, USA. ISBN 1-58113-870-9. doi:10.1145/986537.986635. Online available at <http://doi.acm.org/10.1145/986537.986635> [accessed 2007-09-09].
- [817] Marjan Mernik, Goran Gerlič, Viljem Žumer, and Barrett R. Bryant. Can a parser be generated from examples? In *SAC'03: Proceedings of the 2003 ACM symposium on Applied computing*, pages 1063–1067, 2003, Melbourne, Florida. ACM Press, New York, NY, USA. ISBN 1-58113-624-2. doi:10.1145/952532.952740. Online available at <http://www.cis.uab.edu/softcom/GenParse/sac.pdf> and <http://doi.acm.org/10.1145/952532.952740> [accessed 2007-09-09].
- [818] Matej Črepinček, Marjan Mernik, Faizan Javed, Barrett R. Bryant, and Alan Sprague. Extracting grammar from programs: evolutionary approach. *SIGPLAN Notices*, 40(4):39–46, 2005. ISSN 0362-1340. doi:10.1145/1064165.1064172. Online available at <http://doi.acm.org/10.1145/1064165.1064174> [accessed 2007-09-09].
- [819] Emin Erkan Korkmaz and Göktürk Üçoluk. Genetic Programming for Grammar Induction. In *Genetic and Evolutionary Computation Conference 2001 – Late Breaking Papers*, pages 245–251, 2001. In proceedings [569]. Online available at <http://citeseer.ist.psu.edu/451812.html> and <http://de.scientificcommons.org/464341> [accessed 2007-10-14].
- [820] John R. Koza from the Computer Science Department, Stanford University, Stanford, California 94305 USA. Concept formation and decision tree induction using the genetic programming paradigm. In *Parallel Problem Solving from Nature - Proceedings of 1st Workshop, PPSN 1*, pages 124–128, 1990. In proceedings [592]. Online available at <http://citeseer.ist.psu.edu/61578.html> [accessed 2008-05-29].
- [821] Alex A. Freitas. A Genetic Programming Framework for Two Data Mining Tasks: Classification and Generalized Rule Induction. In *Genetic Programming 1997: Pro-*

- ceedings of the Second Annual Conference GP-97*, pages 96–101, 1997. In proceedings [915]. Online available at <http://citeseer.ist.psu.edu/43454.html> [accessed 2007-09-09].
- [822] Fernando E. B. Otero, Monique M. S. Silva, and Alex A. Freitas. Genetic Programming For Attribute Construction In Data Mining. In *GECCO'02: Proceedings of the Genetic and Evolutionary Computation Conference*, 2002. In proceedings [563] and also [823].
- [823] Fernando E. B. Otero, Monique M. S. Silva, Alex A. Freitas, and J. C. Nievola. Genetic Programming for Attribute Construction in Data Mining. In *Genetic Programming: Proc. 6th European Conference (EuroGP-2003)*, pages 384–393, 2003. In proceedings [906] and also [822]. Online available at http://www.cs.kent.ac.uk/people/staff/aaf/pub_papers.dir/EuroGP-2003-Fernando.pdf [accessed 2007-09-09].
- [824] Celia C. Bojarczuk, Heitor S. Lopes, and Alex A. Freitas. Data mining with constrained-syntax genetic programming: applications to medical data sets. In *Proceedings Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-2001)*, 2001. Online available at <http://citeseer.ist.psu.edu/bojarczuk01data.html> and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/bojarczuk_2001_idamap.html [accessed 2007-09-09].
- [825] Man Leung Wong and Kwong Sak Leung. *Data Mining Using Grammar Based Genetic Programming and Applications*, volume 3 of *Genetic Programming*. Springer, January 2000. ISBN 978-0792377467.
- [826] John Koza, Forrest H. Bennett III, David Andre, and Martin A. Keane. The Design of Analog Circuits by Means of Genetic Programming. In *Evolutionary Design by Computers*, chapter 16, pages 365–385. Morgan Kaufmann, 1999. In collection [599]. Online available at <http://www.genetic-programming.com/jkpdf/edc1999.pdf> [accessed 2007-10-03].
- [827] Tatiana Kalganova. An Extrinsic Function-Level Evolvable Hardware Approach. In *Genetic Programming, Proceedings of EuroGP'2000*, pages 60–75, 2000. In proceedings [909]. Online available at <http://citeseer.ist.psu.edu/kalganova00extrinsic.html> [accessed 2007-09-09].
- [828] John R. Koza, David Andre, Forrest H. Bennett III, and Martin A. Keane. Use of Automatically Defined Functions and Architecture-Altering Operations in Automated Circuit Synthesis Using Genetic Programming. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 132–149, 1996. In proceedings [917]. Online available at <http://citeseer.ist.psu.edu/119355.html> and <http://www.genetic-programming.com/jkpdf/gp1996adfaa.pdf> [accessed 2007-10-03].
- [829] John R. Koza, David Andre, Forrest H. Bennett III, and Martin A. Keane. Evolution of a Low-Distortion, Low-Bias 60 Decibel Op Amp with Good Frequency Generalization using Genetic Programming. In *Late Breaking Papers at the Genetic Programming 1996 Conference*, pages 94–100, 1996. In proceedings [918]. Online available at <http://www.genetic-programming.com/jkpdf/gp1996lbpamplifier.pdf> and <http://citeseer.ist.psu.edu/105348.html> [accessed 2007-10-03].
- [830] John R. Koza, Forrest H. Bennett III, David Andre, and Martin A. Keane. Automatic design of analog electrical circuits using genetic programming. In Hugh Cartwright, editor, *Intelligent Data Analysis in Science*, chapter 8, pages 172–200. Oxford University Press, Oxford, 2000.
- [831] Robert L. Popp, David J. Montana, Richard R. Gassner, Gordon Vidaver, and Suraj Iyer. Automated hardware design using genetic programming, VHDL, and FPGAs. In *IEEE International Conference on Systems, Man, and Cybernetics*, volume 3, pages 2184–2189. IEEE, October 11–14, 1998, San Diego, CA USA. Inspec Accession Number: 6189463.
- [832] Athanasios Tsakonas, Georgios Dounias, Jan Jantzen, Hubertus Axer, Beth Bjerregaard, and Diedrich Graf von Keyserlingk. Evolving rule-based systems in two

- medical domains using genetic programming. *Artificial Intelligence in Medicine*, 32 (3):195–216, 2004. ISSN 0933-3657. doi:10.1016/j.artmed.2004.02.007. Online available at <http://dx.doi.org/10.1016/j.artmed.2004.02.007> [accessed 2007-09-09].
- [833] Markus F. Brameier and Wolfgang Banzhaf. A Comparison of Genetic Programming and Neural Networks in Medical Data Analysis. Technical Report, University of Dortmund, 1998. Reihe Computational Intelligence, Sonderforschungsbereich 531. Online available at <http://citeseer.ist.psu.edu/324837.html> and http://dspace.hrz.uni-dortmund.de:8080/bitstream/2003/5344/2/ci4398_doc.pdf [accessed 2007-09-09].
- [834] Celia C. Bojarczuk, Heitor S. Lopes, and Alex A. Freitas. An innovative application of a constrained-syntax genetic programming system to the problem of predicting survival of patients. In *Genetic Programming: Proc. 6th European Conference (EuroGP-2003)*, pages 11–59, 2003. In proceedings [906]. Online available at http://www.cs.kent.ac.uk/people/staff/aaf/pub_papers.dir/EuroGP-2003-Celia.pdf [accessed 2007-09-09].
- [835] Jin-Hyuk Hong and Sung-Bae Cho. The classification of cancer based on DNA microarray data that uses diverse ensemble genetic programming. *Artificial Intelligence in Medicine*, 36(1):43–58, 2006. Online available at <http://sclab.yonsei.ac.kr/~hjinh/PAPER/IJ2006-1.pdf> and <http://dx.doi.org/10.1016/j.artmed.2005.06.002> [accessed 2007-09-09].
- [836] John R. Koza from the Computer Science Department, Stanford University, Stanford, California 94305, USA. A genetic approach to econometric modeling. In Paul Bourguine and Bernard Walliser, editors, *Economics and Cognitive Science*, pages 57–75. Pergamon Press, Oxford, UK, 1991. Online available at <http://www.genetic-programming.com/jkpdf/ecs1991cecoia1990.pdf> [accessed 2008-05-29]. Submitted under the title “Discovering An Econometric Model By Genetic Breeding Of A Population Of Mathematical Functions” to CECOIA, June 2–6, Paris, France, 1990, The 2nd Conference Economics and Artificial Intelligence, online available at <http://www.genetic-programming.com/jkpdf/cecoia1990.pdf> and <http://citeseer.ist.psu.edu/koza90discovering.html> [accessed 2008-05-29].
- [837] Christopher J. Neely and Paul A. Weller. Predicting exchange rate volatility: genetic programming versus GARCH and RiskMetrics. *Review*, pages 43–54, May 2002. Online available at <http://research.stlouisfed.org/publications/review/02/05/43-54NeelyWeller.pdf> [accessed 2007-09-09].
- [838] Jean-Yves Potvin, Patrick Soriano, and Maxime Vallée. Generating trading rules on the stock markets with genetic programming. *Computers & Operations Research*, 31 (7):1033–1047, June 2004. ISSN 0305-0548. doi:10.1016/S0305-0548(03)00063-7.
- [839] Michael O’Neill, Anthony Brabazon, Conor Ryan, and J. J. Collins. Evolving Market Index Trading Rules Using Grammatical Evolution. In *Proceedings of the EvoWorkshops on Applications of Evolutionary Computing*, pages 343–352, 2001. In proceedings [535]. Online available at <http://citeseer.ist.psu.edu/485032.html> [accessed 2007-09-09].
- [840] Jason D. Lohn, Gregory S. Hornby, and Derek Linden. An evolved antenna for deployment on NASA’s space technology 5 mission. In *Genetic Programming Theory and Practice II*, 2004. In proceedings [922]. Online available at http://ic.arc.nasa.gov/people/hornby/papers/lohn_gptp04.ps.gz [accessed 2007-08-17].
- [841] Koetsu Yamazaki, Sourav Kundu, and Michitomo Hamano. Genetic Programming Based Learning of Control Rules for Variable Geometry Structures. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 412–415, 1998. In proceedings [913]. Online available at <http://citeseer.ist.psu.edu/kundu98genetic.html> [accessed 2007-09-09].
- [842] David Andre, Forrest H. Bennett III, and John R. Koza. Discovery by genetic programming of a cellular automata rule that is better than any known rule for the

- majority classification problem. In *Proceedings of the First Annual Conference Genetic Programming (GP-96)*, pages 3–11, 1996. In proceedings [917] and also [843]. Online available at <http://citeseer.ist.psu.edu/33008.html> [accessed 2007-08-01].
- [843] David Andre, Forrest H. Bennett III, and John R. Koza. Evolution of Intricate Long-Distance Communication Signals in Cellular Automata using Genetic Programming. In *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, volume 1. MIT Press, Cambridge, MA, USA, May 16–18, 1996, Nara, Japan. See also [842]. Online available at <http://www.genetic-programming.com/jkpdf/alife1996gkl.pdf> and <http://citeseer.ist.psu.edu/andre96evolution.html> [accessed 2007-10-03].
- [844] Hugo de Garis. Artificial Embryology: The Genetic Programming of an Artificial Embryo. In *Dynamic, Genetic, and Chaotic Programming*, pages 373–393. John Wiley, 1992. In collection [1912]. Online available at <http://citeseer.ist.psu.edu/99683.html> [accessed 2007-08-01].
- [845] Hugo de Garis from the Universite Libre de Bruxelles, CP 231, Bld. du Triomphe, 1050 Brussels, Belgium, Fax 32-2-659.5767 Phone 32-2-650.5776 Email sgross@ulb.ac.be. Evolving a Replicator: The Genetic Programming of Self Reproduction in Cellular Automata. In *ECAL-93 Self organisation and life: from simple rules to global complexity*, pages 274–284, 1993. In proceedings [1927]. Online available at <http://citeseer.ist.psu.edu/degarris93evolving.html> [accessed 2007-08-01].
- [846] Wolfgang Banzhaf, Peter Nordin, Robert E. Keller, and Frank D. Francone. *Genetic Programming: An Introduction – On the Automatic Evolution of Computer Programs and Its Applications*. Morgan Kaufmann Publishers, first edition, November 30, 1997. ISBN 978-1558605107.
- [847] Sean Luke, Charles Hohn, Jonathan Farris, Gary Jackson, and James Hendler. Co-evolving Soccer Softbot Team Coordination with Genetic Programming. In *Proceedings of the First International Workshop on RoboCup, at the International Joint Conference on Artificial Intelligence*, 1997. Online available at <http://citeseer.ist.psu.edu/3898.html> and <http://www.cs.umd.edu/~seanl/papers/robocup.pdf> [accessed 2007-09-09]. In proceedings [1928, 1929]. See also [848].
- [848] Sean Luke, Charles Hohn, Jonathan Farris, Gary Jackson, and James Hendler. Co-evolving Soccer Softbot Team Coordination with Genetic Programming. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, number 1395 in Lecture Notes in Computer Science (LNCS), subseries Lecture Notes in Artificial Intelligence (LNAI). Springer Berlin/Heidelberg, 1998. ISBN 978-3-540-64473-6. Online available at <http://www.cs.umd.edu/~seanl/papers/robocupc.pdf> [accessed 2007-09-09]. See also [847].
- [849] Sean Luke. Evolving SoccerBots: A Retrospective. In *Proceedings of the 12th Annual Conference of the Japanese Society for Artificial Intelligence (JSAI)*, 1998. Online available at <http://cs.gmu.edu/~sean/papers/robocupShort.pdf> and <http://citeseer.ist.psu.edu/43338.html> [accessed 2007-09-09]. See also [847, 848].
- [850] John R. Koza, Martin A. Keane, Matthew J. Streeter, William Mydlowec, Jessen Yu, and Guido Lanza. *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*. Genetic Programming. Springer, first edition, May 2005. ISBN 978-0387250670.
- [851] John R. Koza and Martin A. Keane. Cart centering and broom balancing by genetically breeding populations of control strategy programs. In *Proceedings of International Joint Conference on Neural Networks*, volume I, pages 198–201, January 15–19, 1990, Washington, USA. Lawrence Erlbaum, Hillsdale, NJ, USA.
- [852] John R. Koza and Martin A. Keane. Genetic breeding of non-linear optimal control strategies for broom balancing. In *In Proceedings of the Ninth International Conference on Analysis and Optimization of Systems*, pages 47–56, June 1990, Antibes, France. Springer-Verlag, Berlin, Germany.

- [853] John R. Koza and James P. Rice. Automatic programming of robots using genetic programming. In *Proceedings of Tenth National Conference on Artificial Intelligence*, pages 194–201, 1992. In proceedings [292]. Online available at <http://citeseer.ist.psu.edu/koza92automatic.html> and <http://www.genetic-programming.com/jkpdf/aaai1992.pdf> [accessed 2007-09-07].
- [854] Hitoshi Iba. Evolving Multiple Agents by Genetic Programming. In *Advances in genetic programming 3*, pages 447–466. MIT Press, Cambridge, MA, USA, 1999. In collection [930]. Online available at <http://www.cs.bham.ac.uk/~wbl/aigp3/ch19.pdf> [accessed 2007-10-03].
- [855] David Andre and Astro Teller. Evolving Team Darwin United. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*, volume 1604 of *Lecture Notes in Artificial Intelligence, subseries of Lecture Notes in Computer Science (LNCS)*, pages 346–351. Springer Verlag, July 1998, Paris, France. ISBN 3-540-66320-7. Published in 1999. Online available at http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/Teller_Astro.ps and <http://citeseer.ist.psu.edu/322830.html> [accessed 2007-10-03].
- [856] Michael O’Neill, J. J. Collins, and Conor Ryan. Automatic Generation of Robot Behaviours using Grammatical Evolution. In *Proceedings of the Fifth International Symposium on Artificial Life and Robotics AROB 2000*, pages 351–354, 2000, Japan.
- [857] Sean Luke and Lee Spector. Evolving Teamwork and Coordination with Genetic Programming. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 150–156, 1996. In proceedings [917]. Online available at <http://citeseer.ist.psu.edu/luke96evolving.html> and <http://www.ifi.uzh.ch/groups/ailab/people/nitschke/refs/Cooperation/cooperation.pdf> [accessed 2008-07-28].
- [858] F. Comellas and G. Giménez from the Department de Matemàtica Aplicada i Telemàtica, Universitat Politècnica de Catalunya, Campus Nord, Edifici C3, 08034 Barcelona, Catalonia, Spain. Genetic Programming to Design Communication Algorithms for Parallel Architectures. *Parallel Processing Letters*, 8(4):549–560, 1998. ISSN 0129-6264. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/Comellas_1998_GPD.html and <http://citeseer.ist.psu.edu/comellas98genetic.html> [accessed 2007-09-14].
- [859] Sérgio Granato de Araújo, Aloysio de Castro Pinto Pedroza, and Antônio Carneiro de Mesquita Filho from the Electrical Engineering Department, Universidade Federal do Rio de Janeiro (UFRJ), Caixa Postal 68504, CEP 21945-970, Rio de Janeiro, RJ, Brazil. Evolutionary Synthesis of Communication Protocols. In *Proceedings of the 10th International Conference on Telecommunications (ICT’03)*, volume 2, pages 986–993, February 23–March 1, 2003, Tahiti, French Polynesia. doi:10.1109/ICTEL.2003.1191573. Online available at <http://www.gta.ufrj.br/ftp/gta/TechReports/AMP03a.pdf> [accessed 2008-06-21]. See also [1646, 1647] for Portuguese versions.
- [860] Thomas Weise from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Genetic Programming for Sensor Networks. Technical Report, University of Kassel, University of Kassel, January 2006. Online available at <http://www.it-weise.de/documents/files/W2006DGPFa.pdf> [accessed 2008-08-12].
- [861] Thomas Weise and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Genetic Programming Techniques for Sensor Networks. In *Proceedings of 5. GI/ITG KuVS Fachgespräch “Drahtlose Sensornetze”*, pages 21–25, July 17–18, 2006, University of Stuttgart, Stuttgart, Germany. Online available at <http://www.it-weise.de/documents/files/W2006DGPFb.pdf> and <http://elib.uni-stuttgart.de/opus/volltexte/2006/2838/> [accessed 2007-11-07].

- [862] Thomas Weise and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. DGPF – An Adaptable Framework for Distributed Multi-Objective Search Algorithms Applied to the Genetic Programming of Sensor Networks. In *Proceedings of the Second International Conference on Bioinspired Optimization Methods and their Application, BIOMA 2006*, pages 157–166, 2006. In proceedings [491]. Online available at <http://www.it-weise.de/documents/files/W2006DGPFc.pdf> [accessed 2008-08-12].
- [863] Thomas Weise, Kurt Geihs, and Philipp Andreas Baer from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Genetic Programming for Proactive Aggregation Protocols. In *Proceedings of the 8th International Conference on Adaptive and Natural Computing Algorithms ICANNGA'07, Part 1*, pages 167–173, 2007. In proceedings [574]. Online available at <http://www.it-weise.de/documents/files/W2007DGPFb.pdf> [accessed 2008-08-12] and http://www.springerlink.com/content/978-3-540-71589-4/?p_o=10 [accessed 2007-08-13].
- [864] Wei Lu and Issa Traore from the Department of Electrical and Computer Engineering, University of Victoria, Victoria, B.C., Canada. Detecting New Forms of Network Intrusions Using Genetic Programming. *Computational Intelligence*, 20(3):475–494, August 2004. Online available at <http://www.isot.ece.uvic.ca/publications/journals/coi-2004.pdf> [accessed 2008-06-11].
- [865] Dong Song, Malcom I. Heywood, and A. Nur Zincir-Heywood. A Linear Genetic Programming Approach to Intrusion Detection. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2003. In proceedings [562]. Online available at <http://users.cs.dal.ca/~mheywood/X-files/Publications/27242325.pdf> [accessed 2008-06-16].
- [866] Dong Song. A Linear Genetic Programming Approach to Intrusion Detection. Master's thesis, Dalhousie University, Halifax, Nova Scotia, Canada, March 2003. Advisors: Malcolm Heywood and Nur Zincir-Heywood. Online available at <http://users.cs.dal.ca/~mheywood/Thesis/DSong.pdf> [accessed 2008-06-16].
- [867] Patrick LaRoche and A. Nur Zincir-Heywood from the Dalhousie University, Halifax, Nova Scotia, Canada. 802.11 Network Intrusion Detection Using Genetic Programming. In *GECCO'05: Proceedings of the 2005 workshops on Genetic and evolutionary computation*, pages 170–171, 2005. doi:10.1145/1102256.1102295. In proceedings [557]. Online available at <http://larocheonline.ca/~plaroche/papers/gecco-laroche-heywood.pdf> and <http://doi.acm.org/10.1145/1102256.1102295> [accessed 2008-06-16].
- [868] Michael Iles and Dwight Lorne Deugo. A Search for Routing Strategies in a Peer-to-Peer Network Using Genetic Programming. In *SRDS'02: Proceedings of the 21st IEEE Symposium on Reliable Distributed Systems (SRDS'02)*, pages 341–346, October 13–16, 2002, Osaka, Japan. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-1659-9. doi:10.1109/RELDIS.2002.1180207. Inspec Accession Number: 7516795.
- [869] I. M. A. Kirkwood, S. H. Shami, and Mark C. Sinclair from the Department of Electronic Systems Engineering, University of Essex, Wivenhoe Park, Colchester, Essex CO4 3SQ, UK. Discovering Simple Fault-Tolerant Routing Rules using Genetic Programming. In *Proceedings of the International Conference on Artificial Neural Nets and Genetic Algorithms*, 1997. ISBN 3-211-83087-1. In proceedings [580]. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/kirkam_1997_dsfttr.html and <http://citeseer.ist.psu.edu/120285.html> [accessed 2008-06-24]. See also [870].
- [870] S. H. Shami, I. M. A. Kirkwood, and Mark C. Sinclair. Evolving simple fault-tolerant routing rules using genetic programming. *Electronics Letters*, 33(17):1440–1441, August 14, 1997. ISSN Print: 0013-5194, Online: 1350-911X. doi:10.1049/el:19970996. See also [869].

- [871] P. Aiyarak, A. S. Saket, and Mark C. Sinclair from the Electronic Systems Engineering, University of Essex, Wivenhoe Park, Colchester, Essex, CO4 3SQ, UK. Genetic Programming Approaches for Minimum Cost Topology Optimisation of Optical Telecommunication Networks. In *Second International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications, GALEZIA*, 1997. In proceedings [744]. Online available at <http://citeseer.ist.psu.edu/30631.html> and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/aiyarak_1997_GPtootn.html [accessed 2008-07-29].
- [872] Mark C. Sinclair from the Electronic Systems Engineering, University of Essex, Wivenhoe Park, Colchester, Essex, CO4 3SQ, UK. Optical Mesh Network Topology Design using Node-Pair Encoding Genetic Programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1192–1197, 1999. In proceedings [572]. Online available at <http://citeseer.ist.psu.edu/177545.html> and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/sinclair_1999_OMNTDNEGP.html [accessed 2008-07-29].
- [873] Mark C. Sinclair. Node-Pair Encoding Genetic Programming for Optical Mesh Network Topology Design. In *Telecommunications Optimization: Heuristic and Adaptive Techniques*, chapter 6, pages 99–114. John Wiley & Sons, Ltd., 2000. In collection [277].
- [874] Forrest H. Bennett III. Emergence of a Multi-Agent Architecture and New Tactics For the Ant Colony Foraging Problem Using Genetic Programming. In *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior: From animals to animats 4*, pages 430–439, 1996. In proceedings [1930].
- [875] Forrest H Bennett III. Automatic Creation of an Efficient Multi-Agent Architecture Using Genetic Programming with Architecture-Altering Operations. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 30–38, 1996. In proceedings [917].
- [876] Mohammad Adil Qureshi. *The Evolution of Agents*. PhD thesis, University College, London, London, UK, July 2001. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/qureshi_thesis.html and <http://citeseer.ist.psu.edu/759376.html> [accessed 2007-09-14].
- [877] Mohammad Adil Qureshi. Evolving Agents. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 369–374, 1996. In proceedings [917] and also [878]. Online available at <http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/AQ.gp96.ps.gz> [accessed 2007-09-17].
- [878] Mohammad Adil Qureshi. Evolving Agents. Research Note RN/96/4, UCL, Gower Street, London, WC1E 6BT, UK, January 1996. Online available at <ftp://ftp.cs.bham.ac.uk/pub/tech-reports/1997/CSRP-97-07.ps.gz> [accessed 2007-09-17]. See also [878].
- [879] Thomas D. Haynes, Roger L. Wainwright, and Sandip Sen from the Department of Mathematical & Computer Sciences, The University of Tulsa. Evolving Cooperation Strategies. Technical Report UTULSA-MCS-94-10, The University of Tulsa, Tulsa, OK, USA, December 16, 1994. See also [1931]. Online available at <http://www.mcs.utulsa.edu/~rogerw/papers/Haynes-icmas95.pdf> and <http://citeseer.ist.psu.edu/1958.html> [accessed 2007-10-03].
- [880] Thomas D. Haynes, Roger L. Wainwright, Sandip Sen, and Dale Schoenefeld. Strongly Typed Genetic Programming in Evolving Cooperation Strategies. In *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 271–278, 1995. In proceedings [747]. Online available at <http://www.mcs.utulsa.edu/~rogerw/papers/Haynes-icga95.pdf> and <http://citeseer.ist.psu.edu/15037.html> [accessed 2007-10-03].
- [881] David Andre. The Automatic Programming of Agents that Learn Mental Models and Create Simple Plans of Action. In *IJCAI-95 Proceedings of the Fourteenth*

- International Joint Conference on Artificial Intelligence*, volume 1, pages 741–747, 1995. MAPMAKER searches for gold. Online available at <http://citeseer.ist.psu.edu/188996.html> and <http://dli.iiit.ac.in/ijcai/IJCAI-95-VOL%201/pdf/097.pdf> [accessed 2007-10-03]. In proceedings [1932].
- [882] Lee Spector and Alan Robinson. Multi-type, Self-adaptive Genetic Programming as an Agent Creation Tool. In *GECCO 2002: Proceedings of the Bird of a Feather Workshops, Genetic and Evolutionary Computation Conference*, pages 73–80, 2002. In proceedings [565]. Online available at <http://citeseer.ist.psu.edu/614297.html> and <http://hampshire.edu/lspector/pubs/ecomas2002-spector-toappear.pdf> [accessed 2007-12-25].
- [883] Peter John Angeline and Jordan B. Pollack from the Laboratory for Artificial Intelligence Research, Computer and Information Science Department, The Ohio State University, Columbus, Ohio 43210, pja@cis.ohio-state.edu, pollack@cis.ohio-state.edu. Coevolving high-level representations. In *Artificial Life III*, pages 55–71, 1992. Online available at <http://demo.cs.brandeis.edu/papers/alife3.pdf> and <http://citeseer.ist.psu.edu/angeline94coevolving.html> [accessed 2007-10-05]. In proceedings [1933].
- [884] Tadahiko Murata and Takashi Nakamura. Genetic network programming with automatically defined groups for assigning proper roles to multiple agents. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1705–1712, 2005. doi:10.1145/1068009.1068294. In proceedings [554]. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2005/docs/p1705.pdf> and <http://doi.acm.org/10.1145/1068009.1068294> [accessed 2008-06-16].
- [885] Hitoshi Iba. Emergent Cooperation for Multiple Agents Using Genetic Programming. In *PPSN IV: Proceedings of the 4th International Conference on Parallel Problem Solving from Nature*, pages 32–41, 1996. doi:10.1007/3-540-61723-X_967. In proceedings [589].
- [886] Hitoshi Iba, Tishihide Nozoe, and Kanji Ueda. Evolving Communicating Agents based on Genetic Programming. In *Proceedings of the 1997 IEEE International Conference on Evolutionary Computation*, pages 297–302, 1997. In proceedings [504]. Online available at <http://lis.epfl.ch/~markus/References/Iba97.pdf> [accessed 2008-07-28].
- [887] Hitoshi Iba from the Machine Inference Section, Electrotechnical Laboratory (ETL), 1-1-4 Umezono, Tsukuba Science City, Ibaraki 305, Japan. Evolutionary learning of communicating agents. *Information Sciences – Informatics and Computer Science: An International Journal*, 108(1-4):181–205, July 1998. ISSN 0020-0255. doi:10.1016/S0020-0255(97)10055-X. Online available at [http://dx.doi.org/10.1016/S0020-0255\(97\)10055-X](http://dx.doi.org/10.1016/S0020-0255(97)10055-X) [accessed 2008-07-28].
- [888] Kenneth J. Mackin and Eiichiro Tazaki from the Department of Control and System Engineering, Toin University of Yokohama, Japan. Emergent agent communication in multi-agent systems using automatically defined function genetic programming (ADF-GP). In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, IEEE SMC'99*, volume 5, pages 138–142, October 12–15, 1999, Tokyo, Japan. ISBN 0780357310, 978-0780357310. doi:10.1109/ICSMC.1999.815536.
- [889] Kenneth J. Mackin and Eiichiro Tazaki from the Department of Control and System Engineering, Toin University of Yokohama, Japan. Unsupervised training of multiobjective agent communication using genetic programming. In *KES*, pages 738–741, 2000. doi:10.1109/KES.2000.884152. In proceedings [347]. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/Mackin00.html> and <http://www.lania.mx/~ccoello/EM00/mackin00.pdf.gz> [accessed 2008-07-28].
- [890] Kenneth J. Mackin and Eiichiro Tazaki from the Department of Control and System Engineering, Toin University of Yokohama, Japan. Multiagent communication

- combining genetic programming and pheromone communication. *Kybernetes*, 31(6): 827–843, 2002. doi:10.1108/03684920210432808. Online available at <http://www.emeraldinsight.com/10.1108/03684920210432808> [accessed 2008-07-28].
- [891] Gregory M. Werner and Michael G. Dyer. Evolution of Communication in Artificial Organisms. In *Artificial Life II*, pages 659–687, 1992. Redwood City, CA. In proceedings [1924]. Online available at <http://www.isrl.uiuc.edu/~amag/langev/paper/werner92evolutionOf.html> [accessed 2008-07-28].
- [892] David Andre. Learning and Upgrading Rules for an OCR System Using Genetic Programming. In *Proceedings of the 1994 IEEE World Congress on Computational Intelligence*, June 27–29, 1994. IEEE Press, Orlando, Florida, USA. Uses GP both to recognise C in various fonts and to maintain manually produced extremely high level code when a new font is added. Online available at <http://citeseer.ist.psu.edu/31976.html> and http://citeseer.ist.psu.edu/cache/papers/cs/802/http:zSzzSzwwww.cs.berkeley.edu/~dandrezSzpapersSzAndre_WCCI_94_OCR_Boundary.pdf/learning-and-upgrading-rules.pdf [accessed 2007-10-03].
- [893] David Andre. Learning and Upgrading Rules for an Optical Character Recognition System Using Genetic Programming. In *Handbook of Evolutionary Computation*. Oxford University Press, 1997. ISBN 0-7503-0392-1. In collection [410].
- [894] Astro Teller and Manuela Veloso. PADO: Learning Tree Structured Algorithms for Orchestration into an Object Recognition System. Technical Report CMU-CS-95-101, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, February 10, 1995. Online available at <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/PADO-Tech-Report.ps> and <http://citeseer.ist.psu.edu/teller95pado.html> [accessed 2008-06-16].
- [895] Astro Teller and Manuela Veloso from the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Algorithm Evolution for Face Recognition: What Makes a Picture Difficult. In *International Conference on Evolutionary Computation*, pages 608–613, 1995. In proceedings [506]. Online available at <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/icecFinal.ps> and <http://citeseer.ist.psu.edu/558916.html> [accessed 2008-06-16].
- [896] Astro Teller and Manuela Veloso. PADO: A New Learning Architecture for Object Recognition. In Katsushi Ikeuchi and Manuela Veloso, editors, *Symbolic Visual Learning*, pages 81–116. Oxford University Press, Inc., New York, NY, USA, May 1, 1996. ISBN 0-19-509870-6, 978-0195098709. Online available at <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/PADO.ps> and <http://citeseer.ist.psu.edu/564427.html> [accessed 2008-06-16].
- [897] John R. Koza and David Andre. Classifying Protein Segments as Transmembrane Domains Using Architecture-Altering Operations in Genetic Programming. In *Advances in Genetic Programming 2*, chapter 8, pages 155–176. MIT Press, 1996. ISBN 0-262-01158-1. In collection [929]. Online available at <http://www.genetic-programming.com/jkpdf/aigp2aatmjk1996.pdf> and <http://citeseer.ist.psu.edu/75759.html> [accessed 2007-10-03].
- [898] John R. Koza and David Andre. Automatic Discovery Using Genetic Programming of an Unknown-Sized Detector of Protein Motifs Containing Repeatedly-used Subexpressions. In *Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications*, pages 89–97, 1995. In proceedings [1934]. Online available at <http://www.genetic-programming.com/jkpdf/ml1995motif.pdf> and <http://citeseer.ist.psu.edu/83080.html> [accessed 2007-10-03].
- [899] John R. Koza and James P. Rice. Genetic Generation of Both the Weights and Architecture for a Neural Network. In *International Joint Conference on Neural Networks, IJCNN-91*, volume II, pages 397–404. IEEE Computer Society Press, July 8–14, 1991, Washington State Convention and Trade Center, Seattle, WA, USA. ISBN 0-7803-

- 0164-1. doi:10.1109/IJCNN.1991.155366. Online available at <http://citeseer.ist.psu.edu/koza91genetic.html> [accessed 2008-05-29].
- [900] Frederic Gruau. Genetic Synthesis of Boolean Neural Networks with a Cell Rewriting Developmental Process. In *Proceedings of the Workshop on Combinations of Genetic Algorithms and Neural Networks (COGANN92)*, pages 55–74, 1992. doi:10.1109/COGANN.1992.273948. In proceedings [1935].
 - [901] Michael O'Neill, Leonardo Vanneschi, Steven Matt Gustafson, Anna Isabel Esparcia Alcázar, Ivanoe De Falco, Antonio Della Cioppa, and Ernesto Tarantino, editors. *Genetic Programming – Proceedings of the 11th European Conference on Genetic Programming, EuroGP 2008*, volume 4971/2008 of *Lecture Notes in Computer Science (LNCS)*, *Theoretical Computer Science and General Issues*, March 26–28 2008, Naples, Italy. Springer-Verlag GmbH, Berlin/Heidelberg, Germany. ISBN 3540786708, 978-3-540-78670-2. doi:10.1007/978-3-540-78671-9.
 - [902] Marc Ebner, Michael O'Neill, Anikó Ekárt, Leonardo Vanneschi, and Anna Isabel Esparcia-Alcázar, editors. *Proceedings of the 10th European Conference on Genetic Programming, EuroGP 2007*, volume 4445/2007 of *Lecture Notes in Computer Science (LNCS)*, April 11-13, 2007, Valencia, Spain. Springer Berlin/Heidelberg. ISBN 978-3-540-71602-0.
 - [903] Pierre Collet, Marco Tomassini, Marc Ebner, Steven Matt Gustafson, and Anikó Ekárt, editors. *Proceedings of the 9th European Conference Genetic Programming, EuroGP 2006*, volume 3905/2006 of *Lecture Notes in Computer Science (LNCS)*, April 10-12, 2006, Budapest, Hungary. Springer Berlin/Heidelberg. ISBN 3-540-33143-3.
 - [904] Maarten Keijzer, Andrea Tettamanzi, Pierre Collet, Jano I. van Hemert, and Marco Tomassini, editors. *Proceedings of the 8th European Conference on Genetic Programming, EuroGP2005*, volume 3447/2005 of *Lecture Notes in Computer Science (LNCS)*, March 30-April 1, 2005, Lausanne, Switzerland. Springer Berlin/Heidelberg. ISBN 3-540-25436-6. See <http://evonet.lri.fr/eurogp2005/> [accessed 2007-09-01].
 - [905] Maarten Keijzer, Una-May O'Reilly, Simon M. Lucas, Ernesto Costa, and Terence Soule, editors. *Proceedings of the 7th European Conference on Genetic Programming, EuroGP2004*, volume 3003/2004 of *Lecture Notes in Computer Science (LNCS)*, April 5-7, 2004, Coimbra, Portugal. Springer Berlin/Heidelberg. ISBN 3-540-21346-5.
 - [906] Conor Ryan, Terence Soule, Maarten Keijzer, Edward P. K. Tsang, Riccardo Poli, and Ernesto Costa, editors. *Proceedings of the 6th European Conference on Genetic Programming, EuroGP 2003*, volume 2610/2003 of *Lecture Notes in Computer Science (LNCS)*, April 14-16, 2003, Essex, UK. Springer Berlin/Heidelberg. ISBN 3-540-00971-X.
 - [907] James A. Foster, Evelynne Lutton, Julian Francis Miller, Conor Ryan, and Andrea Tettamanzi, editors. *Proceedings of the 5th European Conference on Genetic Programming, EuroGP 2002*, volume 2278/2002 of *Lecture Notes in Computer Science (LNCS)*, April 3-5, 2002, Kinsale, Ireland. Springer Berlin/Heidelberg. ISBN 3-540-43378-3.
 - [908] Julian Francis Miller, Marco Tomassini, Pier Luca Lanzi, Conor Ryan, Andrea Tettamanzi, and William B. Langdon, editors. *Proceedings of the 4th European Conference on Genetic Programming, EuroGP 2001*, volume 2038/2001 of *Lecture Notes in Computer Science (LNCS)*, April 18-20, 2001, Lake Como, Italy. Springer Berlin/Heidelberg. ISBN 3-540-41899-7.
 - [909] Riccardo Poli, Wolfgang Banzhaf, William B. Langdon, Julian Francis Miller, Peter Nordin, and Terence C. Fogarty, editors. *Proceedings of the European Conference on Genetic Programming*, volume 1802/2000 of *Lecture Notes in Computer Science (LNCS)*, April 15-16, 2000, Edinburgh, Scotland, UK. Springer Berlin/Heidelberg. ISBN 3-540-67339-3.

- [910] Riccardo Poli, Peter Nordin, William B. Langdon, and Terence C. Fogarty, editors. *Proceedings of the Second European Workshop on Genetic Programming*, volume 1598/1999 of *Lecture Notes in Computer Science (LNCS)*, May 26-27, 1999, Göteborg, Sweden. Springer Berlin/Heidelberg. ISBN 3-540-65899-8.
- [911] Wolfgang Banzhaf, Riccardo Poli, Marc Schoenauer, and Terence C. Fogarty, editors. *Proceedings of the First European Workshop on Genetic Programming, EuroGP'98*, volume 1391/1998 of *Lecture Notes in Computer Science (LNCS)*, April 14-15, 1998, Paris, France. Springer, Berlin/Heidelberg, Germany. ISBN 3-540-64360-5. doi:10.1007/BFb0055923, 978-3-540-64360-9. See also [912].
- [912] Riccardo Poli, William B. Langdon, Marc Schoenauer, Terry Fogarty, and Wolfgang Banzhaf, editors. *Late Breaking Papers at EuroGP'98: The First European Workshop on Genetic Programming*, April 14-15, 1998, Paris, France. Distributed at the workshop. See also [911].
- [913] John R. Koza, Wolfgang Banzhaf, Kumar Chellapilla, Kalyanmoy Deb, Marco Dorigo, David B. Fogel, Max H. Garzon, David E. Goldberg, Hitoshi Iba, and Rick Riolo, editors. *Proceedings of the Third Annual Genetic Programming Conference (GP-98)*, July 22-25, 1998, University of Wisconsin, Madison, Wisconsin, USA. Morgan Kaufmann, Los Altos, CA, USA. ISBN 1-55860-548-7. see <http://www.genetic-programming.org/gp98cfp.html> [accessed 2007-09-01] and [914].
- [914] John R. Koza, editor. *Late Breaking Papers at the Genetic Programming 1998 Conference*, July 22-25, 1998, University of Wisconsin, Madison, Wisconsin, USA. See also [913].
- [915] John R. Koza, Kalyanmoy Deb, Marco Dorigo, David B. Fogel, Max Garzon, Hitoshi Iba, and Rick L. Riolo, editors. *Genetic Programming 1997: Proceedings of the Second Annual Conference GP-97*, July 13-16, 1997, Stanford University, CA, USA. Morgan Kaufmann, San Francisco, CA, USA. See also [916].
- [916] *Late Breaking Papers at the 1997 Genetic Programming Conference*, July 13-16, 1997, Stanford University, CA, USA. Stanford Bookstore. ISBN 0-18-206995-8. See also [915].
- [917] John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors. *Proceedings of the First Annual Conference Genetic Programming (GP-96)*, July 28-31, 1996, Stanford University, CA, USA. MIT Press.
- [918] John R. Koza, editor. *Late Breaking Papers at the First Annual Conference Genetic Programming (GP-96)*, July 28-31, 1996, Stanford University, CA, USA. Stanford Bookstore. ISBN 0-18-201031-7.
- [919] *Genetic Programming Theory and Practice V, Proceedings of the Genetic Programming Theory Practice 2007 Workshop (GPTP-2007)*, Genetic and Evolutionary Computation, May 17-19, 2007, The Center for the Study of Complex Systems (CSCS), University of Michigan, Ann Arbor, Michigan, USA. Springer. See <http://www.cscs.umich.edu/gptp-workshops/gptp2007/> [accessed 2007-09-28].
- [920] Rick Riolo, Terence Soul, and Bill Worzel, editors. *Genetic Programming Theory and Practice IV, Proceedings of the Genetic Programming Theory Practice 2006 Workshop (GPTP-2006)*, Genetic and Evolutionary Computation, May 11-13, 2006, The Center for the Study of Complex Systems (CSCS), University of Michigan, Ann Arbor, Michigan, USA. Springer. ISBN 978-0-387-33375-5. See <http://www.cscs.umich.edu/gptp-workshops/gptp2006/> [accessed 2007-09-28].
- [921] Gwoing Tina Yu, Rick Riolo, and Bill Worzel, editors. *Genetic Programming Theory and Practice III, Proceedings of the Genetic Programming Theory Practice 2005 Workshop (GPTP-2005)*, volume 9 of *Genetic Programming Series*, May 12-14, 2005, The Center for the Study of Complex Systems (CSCS), University of Michigan, Ann Arbor, Michigan, USA. Springer. ISBN 978-0-387-28110-0. See <http://www.cscs.umich.edu/gptp-workshops/gptp2005/> [accessed 2007-09-28].

- [922] Una-May O'Reilly, Gwoing Tina Yu, Rick Riolo, and Bill Worzel, editors. *Genetic Programming Theory and Practice II, Proceedings of the Genetic Programming Theory Practice 2004 Workshop (GPTP-2004)*, volume 8 of *Genetic Programming Series*, May 13-15, 2004, The Center for the Study of Complex Systems (CSCS), University of Michigan, Ann Arbor, Michigan, USA. Springer, New York, USA. ISBN 978-0-387-23253-9, 978-0-387-23254-6. doi:10.1007/b101112. See <http://www.cscs.umich.edu/gptp-workshops/gptp2004/> [accessed 2007-09-28].
- [923] Rick Riolo and Bill Worzel, editors. *Genetic Programming Theory and Practice, Proceedings of the Genetic Programming Theory Practice 2003 Workshop (GPTP-2003)*, Genetic Programming Series, May 15-17, 2003, The Center for the Study of Complex Systems (CSCS), University of Michigan, Ann Arbor, UMichigan, SA. Springer, Kluwer Publishers, Boston, MA. ISBN 1402075812, 978-1402075810. See <http://www.cscs.umich.edu/gptp-workshops/gptp2003/> [accessed 2007-09-28].
- [924] Riccardo Poli, William B. Langdon, and Nicholas Freitag McPhee. *A Field Guide to Genetic Programming*. Published via <http://lulu.com> [accessed 2008-03-26] and freely available at <http://www.gp-field-guide.org.uk> [accessed 2008-03-26], 2008. ISBN 978-1-4092-0073-4. Online available at http://www.lulu.com/items/volume_63/2167000/2167025/2/print/book.pdf and <http://www.gp-field-guide.org.uk/> [accessed 2008-03-27]. (With contributions by John R. Koza).
- [925] John R. Koza. *Genetic Programming II: Automatic Discovery of Reusable Programs: Automatic Discovery of Reusable Programs*. Complex Adaptive Systems. The MIT Press, July 4, 1994. ISBN 0262111896.
- [926] John R. Koza, Forrest H. Bennett III, David Andre, and Martin A. Keane. *Genetic Programming III: Darwinian Invention and Problem Solving*. Morgan Kaufmann, first edition, May 1999. ISBN 978-1558605435.
- [927] William B. Langdon. *Genetic Programming and Data Structures: Genetic Programming + Data Structures = Automatic Programming!* Genetic Programming. Springer, April 30, 1998. ISBN 0792381351.
- [928] Kenneth E. Kinneer, Jr., editor. *Advances in Genetic Programming*, volume 1. MIT Press, Cambridge, MA, USA, April 1994. ISBN 0-262-11188-8.
- [929] Peter J. Angeline and Kenneth E. Kinneer, Jr, editors. *Advances in Genetic Programming*, volume 2. MIT Press, Cambridge, MA, USA, October 1996. ISBN 0-262-01158-1.
- [930] Lee Spector, William B. Langdon, Una-May O'Reilly, and Peter John Angeline, editors. *Advances in Genetic Programming*, volume 3. MIT Press, Cambridge, MA, USA, July 1999. ISBN 0-262-19423-6.
- [931] Markus F. Brameier and Wolfgang Banzhaf. *Linear Genetic Programming*, volume XVI of *Genetic and Evolutionary Computation*. Springer, December 11, 2006. ISBN 0-387-31029-0, 978-0-387-31029-9. Series Editor: David L. Goldberg, John R. Koza.
- [932] Andreas Geyer-Schulz. *Fuzzy Rule-Based Expert Systems and Genetic Machine Learning*, volume 3 of *Studies in Fuzziness and Soft Computing*. Physica-Verlag, 2nd rev edition edition, January 1997. ISBN 978-3790809640. First edition: 1995.
- [933] Lee Spector from the Hampshire College, Amherst, MA, USA. *Automatic Quantum Computer Programming – A Genetic Programming Approach*. Genetic Programming. Kluwer Academic Publishers / Springer Science+Business Media, New York, paperback edition 2007 edition, June 2004. ISBN 0-387-36496-X, 0-387-36791-8, 978-0387-36496-4, 978-0387-36791-0. Series editor: John Koza. Library of Congress Control Number: 2006931640. Partly online available at <http://books.google.de/books?id=mEKfvtxmn9MC> [accessed 2008-05-04].
- [934] Peter Nordin and Wolfgang Banzhaf. Complexity Compression and Evolution. In *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 310–317, 1995. In proceedings [747]. Online available at <http://citeseer.ist.psu.edu/nordin95complexity.html> [accessed 2007-09-07].

- [935] Peter John Angeline. Subtree Crossover: Building Block Engine or Macromutation? In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 9–17, 1997. In proceedings [915].
- [936] Peter John Angeline. Subtree Crossover Causes Bloat. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 745–752. Morgan Kaufmann, 1998. In proceedings [913].
- [937] Frank D. Francone, Markus Conrads, Wolfgang Banzhaf, and Peter Nordin. Homologous Crossover in Genetic Programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1021–1026, 1999. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gecco1999/GP-463.pdf> [accessed 2008-06-16].
- [938] Patrik D’Haeseleer. Context preserving crossover in genetic programming. In *Proceedings of the 1994 IEEE World Congress on Computational Intelligence*, volume 1, pages 256–261, 1994. In proceedings [507]. Online available at <http://citeseer.ist.psu.edu/dhaeseleer94context.html> [accessed 2008-06-15].
- [939] Riccardo Poli and William B. Langdon. A New Schema Theory for Genetic Programming with One-point Crossover and Point Mutation. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 278–285, 1997. In proceedings [915]. Online available at <http://citeseer.ist.psu.edu/327495.html> [accessed 2008-06-14]. See also [940].
- [940] Riccardo Poli and William B. Langdon. A New Schema Theory for Genetic Programming with One-Point Crossover and Point Mutation. Technical Report CSRP-97-3, The University of Birmingham, School of Computer Science, B15 2TT, UK, January 1997. Revised August 1997. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/poli_1997_schemaTR.html and <ftp://ftp.cs.bham.ac.uk/pub/tech-reports/1997/CSRP-97-03.ps.gz> [accessed 2008-06-17]. See also [939].
- [941] David Andre. Evolution of Mapmaking Ability: Strategies for the Evolution of Learning, Planning, and Memory using Genetic Programming. In *Proceedings of the 1994 IEEE World Congress on Computational Intelligence*, volume 1, pages 250–255, June 27–29, 1994. IEEE Press, Orlando, Florida, USA.
- [942] Lee Spector. Evolving Control Structures with Automatically Defined Macros. In E. V. Siegel and John R. Koza, editors, *Working Notes for the AAAI Symposium on Genetic Programming*, pages 99–105, November 10–12, 1995, MIT, Cambridge, MA, USA. AAAI Press, 445 Burgess Drive, Menlo Park, CA 94025, USA. Technical Report FS-95-01. Online available at <http://citeseer.ist.psu.edu/spector95evolving.html> and <http://helios.hampshire.edu/lspector/pubs/ADM-fallsymp-e.pdf> [accessed 2008-06-12].
- [943] Lee Spector. Simultaneous Evolution of Programs and their Control Structures. In *Advances in Genetic Programming 2*, chapter 7, pages 137–154. MIT Press, 1996. In collection [929]. Online available at <http://citeseer.comp.nus.edu.sg/33885.html> and <http://helios.hampshire.edu/lspector/pubs/AiGP2-post-final-e.pdf> [accessed 2008-06-20].
- [944] Wolfgang Banzhaf. Genotype-Phenotype-Mapping and Neutral Variation – A case study in Genetic Programming. In *Parallel Problem Solving from Nature III*, pages 322–332, 1994. In proceedings [590]. Online available at <http://citeseer.ist.psu.edu/banzhaf94genotypephenotypemapping.html> [accessed 2007-09-09].
- [945] Robert E. Keller and Wolfgang Banzhaf from the University of Dortmund, Department of Computer Science, LS11, Germany. Genetic Programming using mutation, reproduction and genotype-phenotype mapping from linear binary genomes into linear LALR(1) phenotypes. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 116–122, 1996. In proceedings [917]. Online available at <http://citeseer.ist.psu.edu/4569.html> [accessed 2007-09-09].

- [946] Robert E. Keller and Wolfgang Banzhaf from the University of Dortmund, Department of Computer Science, LS11, Germany. Genetic Programming using Genotype-Phenotype Mapping from Linear Genomes into Linear Phenotypes. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 116–122, 1996. In proceedings [917]. Online available at <http://citeseer.ist.psu.edu/385878.html> and http://web.cs.mun.ca/~banzhaf/papers/lalr_gp96.ps.gz [accessed 2007-09-09].
- [947] Robert E. Keller and Wolfgang Banzhaf from the University of Dortmund, Department of Computer Science, LS11, Germany. The Evolution of Genetic Code in Genetic Programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1077–1082, 1999. In proceedings [572]. Online available at <http://citeseer.ist.psu.edu/244531.html> and <http://www.cs.mun.ca/~banzhaf/papers/t.pdf> [accessed 2007-09-09].
- [948] Motoo Kimura from the National Institute of Genetics, Mishima, Japan. Evolutionary rate at the molecular level. *Nature*, 217(129):624–626, February 1968.
- [949] Motoo Kimura. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, 1983. ISBN 978-0521317931. reprint edition, 1985, February.
- [950] Richard R. Hudson, Martin Kreitman, and Montserrat Aguade. A Test of Neutral Molecular Evolution Based on Nucleotide Data. *Genetics*, 116(1):153–159, May 1987. Online available at <http://www.genetics.org/cgi/reprint/116/1/153.pdf> [accessed 2007-08-05].
- [951] Cândida Ferreira. Gene Expression Programming: a New Adaptive Algorithm for Solving Problems. *Complex Systems*, 13:87, 2001. Online available at <http://arxiv.org/ftp/cs/papers/0102/0102027.pdf> and <http://www.gene-expression-programming.com/GEPBiblio.asp> [accessed 2007-08-23].
- [952] Cândida Ferreira. Mutation, Transposition, and Recombination: An Analysis of the Evolutionary Dynamics. In H. John Caulfield, Shu-Heng Chen, Heng-Da Cheng, Richard J. Duro, Vasant Honavar, Etienne E. Kerre, Mi Lu, Manuel Grana Roday, Timothy K. Shih, Dan Ventura, Paul P. Wang, and Yuanyuan Yang, editors, *Proceedings of the 6th Joint Conference on Information Science (JCIS)*, pages 614–617. JCIS / Association for Intelligent Machinery, Inc., March 2002, Research Triangle Park, North Carolina, USA. ISBN 0-9707890-1-7. Online available at <http://www.gene-expression-programming.com/GEPBiblio.asp> [accessed 2007-08-23].
- [953] Cândida Ferreira. Analyzing the Founder Effect in Simulated Evolutionary Processes Using Gene Expression Programming. In *Soft Computing Systems - Design, Management and Applications, HIS 2002*, pages 153–162, 2002. In proceedings [308]. Online available at <http://www.gene-expression-programming.com/GEPBiblio.asp> [accessed 2007-08-23].
- [954] Cândida Ferreira. Genetic Representation and Genetic Neutrality in Gene Expression Programming. *Advances in Complex Systems*, 5(4):389–408, 2002. Online available at <http://www.gene-expression-programming.com/GEPBiblio.asp> [accessed 2007-08-23].
- [955] Cândida Ferreira. Function Finding and the Creation of Numerical Constants in Gene Expression Programming. In *7th Online World Conference on Soft Computing in Industrial Applications*, September 2002. Online available at <http://www.gene-expression-programming.com/GEPBiblio.asp> [accessed 2007-08-23].
- [956] Cândida Ferreira. Discovery of the Boolean Functions to the Best Density-Classification Rules Using Gene Expression Programming. In James A. Foster, Evelyne Lutton, Julian Francis Miller, Conor Ryan, and Andrea G. B. Tetta-manzi, editors, *Proceedings of the 5th European Conference on Genetic Programming, EuroGP 2002*, pages 50–59, 2002. ISBN 3-540-43378-3. In proceedings [907]. Online available at <http://www.gene-expression-programming.com/webpapers/ferreira-EuroGP02.pdf> [accessed 2007-09-09].
- [957] Cândida Ferreira. Function Finding and the Creation of Numerical Constants in Gene Expression Programming. In *7th Online World Conference on Soft*

- Computing in Industrial Applications*, September 2002. Online available at http://www.sureserv.com/technic/datum_detail.php?id=466 and <http://www.gene-expression-programming.com/GEPBiblio.asp> [accessed 2007-08-23].
- [958] Heitor S. Lopes and Wagner R. Weinert. EGIPSYS: an Enhanced Gene Expression Programming Approach for Symbolic Regression Problems. *International Journal of Applied Mathematics and Computer Science*, 14(3), 2004. Special Issue: Evolutionary Computation. AMCS Centro Federal de Educacao Tecnologica do Parana / CPGEI Av. 7 de setembro, 3165, 80230-901 Curitiba (PR), Brazil. Online available at <http://matwbn.icm.edu.pl/ksiazki/amc/amc14/amc1437.pdf> [accessed 2007-08-23].
- [959] Chi Zhou, Peter C. Nelson, Weimin Xiao, and Thomas M. Tirpak. Discovery of Classification Rules by Using Gene Expression Programming. In *Proceedings of the International Conference on Artificial Intelligence (IC-AI'02)*, pages 1355–1361, June 2002. Las Vegas, U.S.A.
- [960] Jie Zuo, Changjie Tang, and Tianqing Zhang. Mining Predicate Association Rule by Gene Expression Programming. In Xiaofeng Meng, Jianwen Su, and Yujun Wang, editors, *WAIM'02: Proceedings of the Third International Conference on Advances in Web-Age Information Management*, volume 2419/2002 of *Lecture Notes in Computer Science (LNCS)*, pages 281–294. Springer-Verlag, August 11–13, 2002, Beijing, China. ISBN 3-540-44045-3. Computer Department, Sichuan University China.
- [961] Chi Zhou, Weimin Xiao, Thomas M. Tirpak, and Peter C. Nelson. Evolving accurate and compact classification rules with gene expression programming. *IEEE Transactions on Evolutionary Computation*, 7(6):519–531, December 2003. ISSN 1089-778X.
- [962] M. H. Marghny and I. E. El-Semman from the Department of Computer Science, Faculty of Computers and Information, Assuit University, Egypt. Extracting fuzzy classification rules with gene expression programming. *ICGST International Journal on Artificial Intelligence and Machine Learning, AIML*, (Special Issue on AI & Specific Applications), 2006. AIML 05 Conference, 19-21 December 2005, CICC, Cairo, Egypt. Online available at <http://www.icgst.com/AIML05/papers/P1120535114.pdf> [accessed 2007-08-23].
- [963] Hongqing Cao, Jingxian Yu, and Lishan Kang. An evolutionary approach for modeling the equivalent circuit for electrochemical impedance spectroscopy. In *Proceedings of the 2003 Congress on Evolutionary Computation CEC2003*, pages 1819–1825, 2003. In proceedings [498].
- [964] Candida Ferreira. Designing Neural Networks Using Gene Expression Programming. In Ajith Abraham and Mario Köppen, editors, *9th Online World Conference on Soft Computing in Industrial Applications*, September 2004. Online available at <http://www.gene-expression-programming.com/GEPBiblio.asp> [accessed 2007-08-23].
- [965] Kangshun Li, Yuanxiang Li, Haifang Mo, and Zhangxin Chen. A New Algorithm of Evolving Artificial Neural Networks via Gene Expression Programming. *Journal of the Korean Society for Industrial and Applied Mathematics*, 9(2), 2005.
- [966] Liliana Teodorescu. High energy physics data analysis with gene expression programming. In *Nuclear Science Symposium Conference Record*, volume 1, pages 143–147. IEEE, October 23–29, 2005. doi:10.1109/NSSMIC.2005.1596225. INSPEC Accession Number:8976991.
- [967] M. E. Keskin and Özlem Terzi. Modeling Water Temperature Using Gene Expression Programming. In *Proceedings of the 14th Turkish Symposium on Artificial Intelligence and Neural Networks, TAINN 2005*, pages 280–285, 2005, Izmir, Turkey.
- [968] Abdulkadir Çevik from the Department of Civil Engineering, University Of Gaziantep, Turkey. A new formulation for web crippling strength of cold-formed steel sheeting using genetic programming. *Journal of Constructional Steel Research*, 63(7):867–883, July 2007. doi:10.1016/j.jcsr.2006.08.012.
- [969] Sean Luke and Lee Spector. Evolving Graphs and Networks with Edge Encoding: A Preliminary Report. In *Late Breaking Papers at the First Annual Conference Ge-*

- netic Programming (GP-96)*, 1996. In proceedings [918]. Online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.25.9484> and <http://cs.gmu.edu/~sean/papers/graph-paper.pdf> [accessed 2008-08-02].
- [970] David J. Montana. Strongly Typed Genetic Programming. BBN Technical Report #7866, Bolt Beranek and Newman Inc. (BBN), 70 Fawcett Street, Cambridge, MA 02138, dmontana@bbn.com, May 7, 1993. Online available at <http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/ftp.io.com/papers/stgp.ps.Z> [accessed 2007-10-04]. Superseded by [971].
- [971] David J. Montana. Strongly Typed Genetic Programming. BBN Technical Report #7866, Bolt Beranek and Newman Inc. (BBN), 70 Fawcett Street, Cambridge, MA 02138, dmontana@bbn.com, March 25, 1994. Online available at <http://citeseer.ist.psu.edu/345471.html> and <http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/ftp.io.com/papers/stgp2.ps.Z> [accessed 2007-10-04]. Supersedes [970] and is itself superseded by [972].
- [972] David J. Montana. Strongly Typed Genetic Programming. *Evolutionary Computation*, 3(2):199–230, 1995. Online available at <http://vishnu.bbn.com/papers/stgp.pdf> [accessed 2007-10-04] (November 20, 2002 edition).
- [973] Thomas D. Haynes, Dale A. Schoenefeld, and Roger L. Wainwright. Type Inheritance in Strongly Typed Genetic Programming. In *Advances in Genetic Programming 2*, chapter 18, pages 359–376. MIT Press, 1996. In collection [929]. Online available at <http://citeseer.ist.psu.edu/haynes96type.html> [accessed 2007-10-04].
- [974] Hendrik James Antonisse. A Grammar-Based Genetic Algorithm. In *Proceedings of Foundations of Genetic Algorithms FOGA-90*, pages 193–204, 1990. In proceedings [188].
- [975] Pawel A. Stefanski. Genetic Programming Using Abstract Syntax Trees, 1993. Notes from Genetic Programming Workshop at ICGA-93 [748]. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/icga93-gp-stefanski.html> [accessed 2007-08-15] (The file is actually a zip-archive.).
- [976] Gerald P. Roston. *A Genetic Methodology for Configuration Design*. PhD thesis, Department of Mechanical Engineering of Carnegie Mellon University, Pittsburgh, PA 15213-3891, USA, December 1994. Advisors: Rober Sturges, Jr. and William “Red” Wittaker. Online available at <http://citeseer.ist.psu.edu/213003.html> and http://www.ri.cmu.edu/pubs/pub_3335.html [accessed 2007-08-15].
- [977] Jun’ichi Mizoguchi, Hitoshi Hemmi, and Katsunori Shimohara. Production Genetic Algorithms for Automated Hardware Design Through an Evolutionary Process. In *IEEE World Congress on Computational Intelligence, Proceedings of the First IEEE Conference on Evolutionary Computation*, volume 2, pages 661–664. IEEE Press, June 27–29, 1994. doi:10.1109/ICEC.1994.349980.
- [978] Conor Ryan, Michael O’Neill, and J. J. Collins. Grammatical Evolution: Solving Trigonometric Identities. In *Proceedings of Mendel 1998: 4th International Mendel Conference on Genetic Algorithms, Optimisation Problems, Fuzzy Logic, Neural Networks, Rough Sets*, pages 111–119, 1998. In proceedings [381]. Online available at <http://citeseer.ist.psu.edu/360445.html> [accessed 2007-11-12].
- [979] Man Leung Wong and Kwong Sak Leung. Combining Genetic Programming and Inductive Logic Programming using Logic Grammars. In *IEEE Conference on Evolutionary Computation*, volume 2, pages 733–736, 1995. In proceedings [506]. Online available at <http://cptra.ln.edu.hk/staffProfile/mlwongPub.htm> and <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/ManLeungWong.html> [accessed 2007-08-15].
- [980] Man Leung Wong and Kwong Sak Leung. Learning First-order Relations from Noisy Databases using Genetic Algorithms. In *Proceedings of the Second Singapore International Conference on Intelligent Systems*, pages 159–164, 1994. Online available at <http://cptra.ln.edu.hk/staffProfile/mlwongPub.htm> and <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/ManLeungWong.html> [accessed 2007-08-15].

- [981] Man Leung Wong and Kwong Sak Leung. An adaptive Inductive Logic Programming system using Genetic Programming. In *Evolutionary Programming IV Proceedings of the Fourth Annual Conference on Evolutionary Programming*, pages 737–752, 1995. In proceedings [1189], Online available at <http://cptra.ln.edu.hk/staffProfile/mlwongPub.htm> and <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/ManLeungWong.html> [accessed 2007-08-15].
- [982] Man Leung Wong and Kwong Sak Leung. Inducing logic programs with genetic algorithms: the Genetic Logic Programming System Genetic Logic Programming and Applications. *IEEE Expert*, 10(5):68–76, October 1995. doi:10.1109/64.464935. IEEE Expert Special Track on Evolutionary Programming (Peter John Angeline ed.). Online available at <http://cptra.ln.edu.hk/staffProfile/mlwongPub.htm> and <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/ManLeungWong.html> [accessed 2007-08-15].
- [983] Man Leung Wong and Kwong Sak Leung. Applying logic grammars to induce sub-functions in genetic programming. In *Proceedings of IEEE International Conference on Evolutionary Computation*, volume 2, 1995. In proceedings [506]. Online available at <http://cptra.ln.edu.hk/staffProfile/mlwongPub.htm> and <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/ManLeungWong.html> [accessed 2007-08-15].
- [984] Man Leung Wong and Kwong Sak Leung. Evolutionary Program Induction Directed by Logic Grammars. *Evolutionary Computation*, 5(2):143–180, summer 1997. Special Issue: Trends in Evolutionary Methods for Program Induction. Online available at <http://cptra.ln.edu.hk/staffProfile/mlwongPub.htm> and <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/ManLeungWong.html> [accessed 2007-08-15].
- [985] Peter Alexander Whigham. Context-Free Grammar and Genetic Programming. Technical Report Technical Report CS20/94, Department of Computer Science, Australian Defence Force Academy, University of New South Wales, Canberra ACT 2600, Australia, 1994.
- [986] Peter Alexander Whigham. Grammatically-based Genetic Programming. In *Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications*, pages 33–41, 1995. In proceedings [1934]. Online available at <http://citeseer.ist.psu.edu/whigham95grammaticallybased.html> [accessed 2007-08-15].
- [987] P. A. Whigham. Inductive Bias and Genetic Programming. In *First International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications, GALEZIA*, pages 461–466, 1995. In proceedings [745]. Online available at <http://citeseer.ist.psu.edu/343730.html> [accessed 2008-08-15].
- [988] Peter Alexander Whigham. Search Bias, Language Bias, and Genetic Programming. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 230–237, 1996. In proceedings [917]. Online available at <http://citeseer.ist.psu.edu/whigham96search.html> and ftp://www.cs.adfa.edu.au/pub/xin/whigham_gp96.ps.gz [accessed 2007-09-09].
- [989] Peter Alexander Whigham. *Grammatical bias for evolutionary learning*. PhD thesis, School of Computer Science, University College, University of New South Wales, Australian Defence Force Academy, Canberra, New South Wales, Australia, October 14, 1996. Order Number: AAI0597571.
- [990] Walter Bohm and Andreas Geyer-Schulz. Exact Uniform Initialization for Genetic Programming. In *Foundations of Genetic Algorithms IV*, pages 379–407, 1996. In proceedings [736]. k-bounded context-free languages.
- [991] Helmut Hörner. A C++ Class Library for GP: Vienna University of Economics Genetic Programming Kernel (Release 1.0, operating instructions). Technical Report, Vienna University of Economics, May 29, 1996. Online available at <http://citeseer.ist.psu.edu/253491.html> [accessed 2007-09-09].
- [992] Norman R. Paterson and Mike Livesey. Distinguishing Genotype and Phenotype in Genetic Programming. In *Late Breaking Papers at the Genetic Programming*

- 1996 Conference, pages 141–150, 1996. In proceedings [918]. Online available at <http://citeseer.ist.psu.edu/82479.html> and <ftp://ftp.dcs.st-and.ac.uk/pub/norman/GADS.ps.gz> [accessed 2007-09-07].
- [993] Norman R. Paterson and Mike Livesey. Evolving caching algorithms in C by GP. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 262–267, 1997. In proceedings [915].
- [994] Michael O'Neill. Grammatical Evolution. In *Proceedings of the Fifth Research Conference of the Department of Computer Science and Information Systems, University of Limerick*, September 1998.
- [995] Conor Ryan, J. J. Collins, and Michael O'Neill. Grammatical Evolution: Evolving Programs for an Arbitrary Language. In *Proceedings of the First European Workshop on Genetic Programming*, pages 83–95, 1998. In proceedings [911]. Online available at <http://www.grammatical-evolution.org/papers/eurogp98.ps> and <http://citeseer.ist.psu.edu/ryan98grammatical.html> [accessed 2007-09-09].
- [996] Conor Ryan and Michael O'Neill. Grammatical Evolution: A Steady State approach. In *Late Breaking Papers at the Genetic Programming 1998 Conference*, 1998. In proceedings [914]. Online available at <http://citeseer.ist.psu.edu/260828.html> and <http://www.grammatical-evolution.org/papers/gp98/index.html> [accessed 2007-09-09].
- [997] Michael O'Neill, Finbar Leahy, and Anthony Brabazon. Grammatical Swarm: A variable-length Particle Swarm Algorithm. In *Swarm Intelligent Systems*, chapter 5. Springer, 2006. In collection [1437].
- [998] Michael O'Neill and Anthony Brabazon. Grammatical Swarm: The generation of programs by social programming. *Natural Computing: an international journal*, 5(4):443–462, November 2006. ISSN 1567-7818 (Print) 1572-9796 (Online). doi:10.1007/s11047-006-9007-7. Online available at <http://dx.doi.org/10.1007/s11047-006-9007-7> and <http://www.springerlink.com/content/p3t1923gr7725583/fulltext.pdf> [accessed 2007-09-09].
- [999] Michael O'Neill and Anthony Brabazon. Grammatical Differential Evolution. In Hamid R. Arabnia, editor, *Proceedings of the 2006 International Conference on Artificial Intelligence, ICAI 2006*, volume 1, pages 231–236. CSREA Press, June 26–29, 2006, Las Vegas, Nevada, USA. ISBN 1-932415-96-3. Online available at <http://ww1.ucmss.com/books/LFS/CSREA2006/ICA4864.pdf> [accessed 2007-09-09].
- [1000] Conor Ryan from the Biocomputing and Developmental Systems Group, Department of Computer Science and Information Systems, University of Limerick. Grammatical Evolution Tutorial. In *Proceedings of Genetic and Evolutionary Computation Conference, GECCO 2006*, 2006. In proceedings [553]. Online available at <http://www.grammaticalevolution.org/tutorial.pdf> [accessed 2007-09-09].
- [1001] Michael O'Neill and Conor Ryan. Evolving Multi-line Compilable C Programs. In *Proceedings of the Second European Workshop on Genetic Programming*, pages 83–92, 1999. In proceedings [910]. Online available at <http://citeseer.ist.psu.edu/278127.html> and <http://www.grammatical-evolution.org/papers/eurogp99.ps.gz> [accessed 2007-09-09].
- [1002] Michael O'Neill. *Automatic Programming in an Arbitrary Language: Evolving Programs with Grammatical Evolution*. PhD thesis, University of Limerick, 2001.
- [1003] Michael O'Neill, J. J. Collins, and Conor Ryan. Automatic Programming of Robots. In *Proceedings of Artificial Intelligence and Cognitive Science AICS 2000*, 2000.
- [1004] Anthony Brabazon, Michael O'Neill, Robin Matthews, and Conor Ryan. Grammatical Evolution And Corporate Failure Prediction. In *GECCO'02: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1011–1018, 2002. In proceedings [563]. Online available at <http://business.kingston.ac.uk/research/intbus/paper4.pdf> and <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2002/RWA145.pdf> [accessed 2007-09-09].

- [1005] Anthony Brabazon and Michael O'Neill. Evolving financial models using grammatical evolution. In *Proceedings of The Annual Conference of the South Eastern Accounting Group (SEAG) 2003*, September 8, 2003, London Metropolitan University, London.
- [1006] Anthony Brabazon and Michael O'Neill. A Grammar Model for Foreign Exchange Trading. In H. R. Arabnia et al., editor, *Proceedings of the International Conference on Artificial Intelligence*, volume II, pages 492–498. CSREA Press, June 2003. ISBN 1-932415-13-0.
- [1007] Michael O'Neill and Conor Ryan from the Biocomputing and Developmental Systems Dept. Of Computer Science & Information Systems University of Limerick, Ireland. Grammatical Evolution by Grammatical Evolution: The Evolution of Grammar and Genetic Code. In *Proceedings of 7th European Conference on Genetic Programming, EuroGP2004*, pages 138–149, 2004. In proceedings [905]. Online available at <http://ncra.ucd.ie/papers/eurogp2004.pdf> [accessed 2007-08-28].
- [1008] Michael O'Neill and Conor Ryan, editors. *The 3rd Grammatical Evolution Workshop (GEWS 2004)*, June 26, 2004, Seattle, WA, USA. Part of GECCO, see [558, 559] and also <http://www.grammatical-evolution.com/gews2004/> [accessed 2007-09-10].
- [1009] Michael O'Neill and Conor Ryan, editors. *The 2nd Grammatical Evolution Workshop (GEWS 2003)*, July 2003, Chicago, IL, USA. Part of GECCO, see [561, 562] and also <http://www.grammatical-evolution.com/gews2003/> [accessed 2007-09-10].
- [1010] John N. Shutt. Recursive Adaptable Grammars. Master's thesis, Computer Science Department, Worcester Polytechnic Institute, Worcester Massachusetts, August 10, 1993. Approved by Roy S. Rubinstein and Robert E. Kinicki. Online available at <http://libra.msra.cn/paperdetail.aspx?id=256609> and <http://en.scientificcommons.org/234810> [accessed 2007-08-17].
- [1011] Ronald Morrison. *On the Development of Algol*. PhD thesis, Department of Computational Science, University of St Andrews, December 1979. Online available at <http://www-old.cs.st-andrews.ac.uk/research/publications/Mor79a.php> [accessed 2007-07-10].
- [1012] Sean Luke, Liviu Panait, Gabriel Balan, Sean Paus, Zbigniew Skolicki, Jeff Bassett, Robert Hubley, and Alexander Chircop. ECJ: A Java-based evolutionary computation research system, 2006. In 2007, ECJ reached version 16. For more information <http://cs.gmu.edu/~eclab/projects/ecj/> [accessed 2007-07-10].
- [1013] Henning Christiansen. Syntax, semantics, and implementation strategies for programming languages with powerful abstraction mechanisms. In *Proceedings of 18th Hawaii International Conference on System Sciences*, volume 2, pages 57–66, 1985. Also Datalogiske skrifter 1, Department of Computer Science, Roskilde University, 1985.
- [1014] Marina de la Cruz Echeandía, Alfonso Ortega de la Puente, and Manuel Alfonso. Attribute Grammar Evolution. In José Mira and José R. Álvarez, editors, *Artificial Intelligence and Knowledge Engineering Applications: A Bioinspired Approach, Part II*, volume 3562/2005 of *Lecture Notes in Computer Science (LNCS)*, pages 182–191. Springer Berlin / Heidelberg, August 2005. ISBN 978-3-540-26319-7. doi:10.1007/b137296. Online available at <http://arantxa.ii.uam.es/~alfonsec/docs/confint/iwinac05.pdf> [accessed 2007-09-09].
- [1015] Alfonso Ortega de la Puente, Marina de la Cruz Echeandía, and Manuel Alfonso from the Departamento de Ingeniería Informática de la Escuela Politécnica Superior de la Universidad Autónoma de Madrid, España. Christiansen Grammar Evolution: Grammatical Evolution With Semantics. *IEEE Transactions on Evolutionary Computation*, 11(1):77–90, February 2007. ISSN 1089-778X. Online available at <http://arantxa.ii.uam.es/~alfonsec/artint/ieeetec.pdf> [accessed 2007-09-09].
- [1016] Xuan Hoai Nguyen and Robert Ian McKay. A Framework for Tree-adjunct Grammar Guided Genetic Programming. In *Proceedings of the Post-graduate ADFA Conference on Computer Science (PACCS'01)*, volume 1 of *ADFA Monographs in Computer*

- Science Series*, pages 93–100, July 14, 2001, Canberra ACT, Australia. ISBN 0-7317-0507-6. Online available at <http://sc.snu.ac.kr/PAPERS/TAG3P.pdf> [accessed 2007-09-09].
- [1017] Xuan Hoai Nguyen, Robert Ian McKay, D. L. Essam, and R. Chau. Solving the symbolic regression problem with tree-adjunct grammar guided genetic programming: the comparative results. In *CEC'02: Proceedings of the Congress on Evolutionary Computation 2002*, pages 1326–1331, 2002. In proceedings [499]. Online available at <http://sc.snu.ac.kr/PAPERS/MCEC2002.pdf> [accessed 2007-09-09].
- [1018] Xuan Hoai Nguyen, Robert Ian McKay, and D. L. Essam. Some Experimental Results with Tree Adjunct Grammar Guided Genetic Programming. In *EuroGP'02: Proceedings of the 5th European Conference on Genetic Programming*, pages 229–238, 2002. In proceedings [907]. Online available at <http://sc.snu.ac.kr/PAPERS/TAGGGP.pdf> [accessed 2007-09-09].
- [1019] Xuan Hoai Nguyen, Robert Ian McKay, and D. L. Essam. Can Tree Adjunct Grammar Guided Genetic Programming be Good at Finding a Needle in a Haystack? – A Case Study. In *Proceedings of IEEE International Conference on Communications, Circuits and Systems and West Sino Expositions*, volume 2, pages 1113–1117. IEEE Press, June 29–July 1, 2002. Online available at <http://sc.snu.ac.kr/PAPERS/hoaietal.pdf> [accessed 2007-09-09].
- [1020] Xuan Hoai Nguyen, Robert Ian McKay, and H. A. Abbass. Tree Adjoining Grammars, Language Bias, and Genetic Programming. In *Genetic Programming: 6th European Conference*, pages 157–183, 2003. In proceedings [906]. Online available at <http://sc.snu.ac.kr/PAPERS/eurogp03.pdf> [accessed 2007-09-10].
- [1021] David Jeremy Weir. *Characterizing Mildly Context-Sensitive Grammar Formalisms*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 1988. Supervisor: Aravind K. Joshi, Order Number:AAI8908403.
- [1022] Aravind K. Joshi and Yves Schabes. Tree-Adjoining Grammars. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, pages 69–124. Springer, Berlin, New York, 1997. Online available at <http://citeseer.ist.psu.edu/joshi97treeadjoining.html> [accessed 2007-09-15].
- [1023] Xuan Hoai Nguyen. Solving Trigonometric Identities with Tree Adjunct Grammar Guided Genetic Programming. In *Hybrid Information Systems, Proceedings of First International Workshop on Hybrid Intelligent Systems*, pages 339–351, 2001. In proceedings [309]. See also [1024].
- [1024] Xuan Hoai Nguyen, Robert Ian McKay, and D. L. Essam. Finding Trigonometric Identities with Tree Adjunct Grammar Guided Genetic Programming. In Ajith Abraham, Lakhmi C. Jain, and Berend van der Zwaag, editors, *Innovations in Intelligent Systems and Applications*, volume 140 of *Springer Studies in Fuzziness and Soft Computing*, pages 221–236. Springer-Verlag, January 2004. ISBN 978-3-540-20265-3. See also [1023]. Online available at <http://sc.snu.ac.kr/PAPERS/trigonometry.pdf> [accessed 2007-09-09].
- [1025] Peter Nordin. A compiling genetic programming system that directly manipulates the machine code. In *Advances in genetic programming 1*, chapter 14, pages 311–331. The MIT Press, 1994. In collection [928]. Machine code GP Sun Spark and i868.
- [1026] Markus F. Brameier and Wolfgang Banzhaf. A Comparison of Linear Genetic Programming and Neural Networks in Medical Data Mining. *IEEE Transactions on Evolutionary Computation*, 5(1):17–26, 2001. Online available at <http://citeseer.ist.psu.edu/brameier00comparison.html> and http://web.cs.mun.ca/~banzhaf/papers/ieee_taec.pdf [accessed 2007-09-09].
- [1027] Markus F. Brameier and Wolfgang Banzhaf from the Department of Computer Science, University of Dortmund, 44221 Dortmund, Germany. Evolving Teams of Predictors with Linear Genetic Programming. *Genetic Programming and Evolvable Machines*, 2(4):381–407, December 2001. ISSN 1389-2576 (Print) 1573-7632 (Online).

- doi:10.1023/A:1012978805372. Online available at <http://dx.doi.org/10.1023/A:1012978805372> [accessed 2007-09-09].
- [1028] Markus F. Brameier and Wolfgang Banzhaf. Neutral Variations Cause Bloat in Linear GP. In *Proceedings of 6th European Conference on Genetic Programming, EuroGP*, pages 997–1039, 2003. In proceedings [906].
- [1029] Markus F. Brameier. *On Linear Genetic Programming*. PhD thesis, Fachbereich Informatik, Universität Dortmund, February 2004, Dortmund. Day of Submission: 2003-05-28, Committee: Wolfgang Banzhaf and Martin Riedmiller and Peter Nordin. Online available at <https://eldorado.uni-dortmund.de/handle/2003/20098> and <http://hdl.handle.net/2003/20098> [accessed 2007-08-17].
- [1030] Peter Nordin, Frank Francone, and Wolfgang Banzhaf. Explicitly Defined Introns and Destructive Crossover in Genetic Programming. In *Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications*, pages 6–22, 1995. In proceedings [1934] and also [1118]. Online available at <http://citeseer.ist.psu.edu/nordin95explicitly.html> and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/nordin_1995_introns.html [accessed 2007-09-07].
- [1031] Malcom I. Heywood and A. Nur Zincir-Heywood. Dynamic Page-Based Linear Genetic Programming. *IEEE Transactions on Systems, Man and Cybernetics – Part B: Cybernetics*, 32(3):380–388, 2002. ISSN 1083-4419.
- [1032] Astro Teller. Evolving Programmers: The Co-evolution of Intelligent Recombination Operators. In *Advances in Genetic Programming 2*, chapter 3, pages 45–68. MIT Press, 1996. In collection [929]. Online available at <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/AiGPII.ps> and <http://www.cs.cmu.edu/afs/cs/usr/astro/mosaic/chapterII/chapterII.html> [accessed 2008-06-16].
- [1033] Riccardo Poli from the School of Computer Science, The University of Birmingham, Birmingham B15 2TT, United Kingdom. Parallel Distributed Genetic Programming. Technical Report CSRP-96-15, School of Computer Science, The University of Birmingham, Birmingham B15 2TT, United Kingdom, September 1996. See [1036]. Online available at <http://citeseer.ist.psu.edu/86223.html> [accessed 2007-11-04].
- [1034] Riccardo Poli from the School of Computer Science, The University of Birmingham, Birmingham B15 2TT, United Kingdom. Some Steps Towards a Form of Parallel Distributed Genetic Programming. In *The 1st Online Workshop on Soft Computing (WSC1)*. Nagoya University, Japan, August 19-30, 1996. Originally published as technical report, CSRP-96-14, University of Birmingham, School of Computer Science, 1996. Online available at <http://cswww.essex.ac.uk/staff/poli/papers/Poli-WSC1-1996.pdf> and <http://citeseer.ist.psu.edu/156274.html> [accessed 2007-11-04].
- [1035] Riccardo Poli from the School of Computer Science, The University of Birmingham, Birmingham B15 2TT, United Kingdom. Evolution of Graph-like Programs with Parallel Distributed Genetic Programming. In *Genetic Algorithms: Proceedings of the Seventh International Conference*, pages 346–353, 1997. In proceedings [746]. Online available at <http://citeseer.ist.psu.edu/578015.html> and <http://cswww.essex.ac.uk/staff/poli/papers/Poli-ICGA1997-PDGP.pdf> [accessed 2007-11-04].
- [1036] Riccardo Poli from the School of Computer Science, The University of Birmingham, Birmingham B15 2TT, United Kingdom. Parallel Distributed Genetic Programming. In *New Ideas in Optimization*, pages 403–432. McGraw-Hill Education, 1999. In collection [400], see also [1033]. Online available at <http://cswww.essex.ac.uk/staff/rpoli/papers/Poli-NIO-1999-PDGP.pdf> [accessed 2007-11-04].
- [1037] Riccardo Poli from the School of Computer Science, The University of Birmingham, Birmingham B15 2TT, United Kingdom. Discovery of Symbolic, Neuro-Symbolic and Neural Networks with Parallel Distributed Genetic Programming. In *3rd International Conference on Artificial Neural Networks and Genetic Algorithms, ICANNGA'97*, 1997. In proceedings [580]. Online available at <http://cswww.essex.ac.uk>.

- uk/staff/rpoli/papers/Poli-ICANNGA1997.pdf and <http://citeseer.ist.psu.edu/poli97discovery.html> [accessed 2007-11-04].
- [1038] Riccardo Poli. Evolution of Recursive Transition Networks for Natural Language Recognition with Parallel Distributed Genetic Programming. Technical Report CSRP-96-19, School of Computer Science, The University of Birmingham, Birmingham B15 2TT, United Kingdom, December 1996. Online available at <http://citeseer.ist.psu.edu/90834.html> [accessed 2007-11-04]. See also [1039].
 - [1039] Riccardo Poli. Evolution of Recursive Transition Networks for Natural Language Recognition with Parallel Distributed Genetic Programming. In David Corne and Jonathan L. Shapiro, editors, *Proceedings of the Workshop on Artificial Intelligence and Simulation of Behaviour (AISB) International Workshop on Evolutionary Computing*, volume 1305 of *Lecture Notes in Computer Science (LNCS)*, pages 163–177. Springer-Verlag, April 7-8, 1997, Manchester, UK. ISBN 3-540-63476-2. Online available at <http://cswww.essex.ac.uk/staff/rpoli/papers/Poli-AISB-1997.pdf> and <http://citeseer.ist.psu.edu/355686.html> [accessed 2007-11-04]. See also [1038].
 - [1040] Hironobu Katagiri, Kotaro Hirasawa, Jinglu Hu, and Junichi Murata. Network Structure Oriented Evolutionary Model– Genetic Network Programming – and its Comparison with Genetic Programming. In *Genetic and Evolutionary Computation Conference Late Breaking Papers*, pages 219–226, 2001. In proceedings [569]. See also [1041, 1042].
 - [1041] Hironobu Katagiri, Kotaro Hirasawa, Jinglu Hu, and Junichi Murata. Network Structure Oriented Evolutionary Model – Genetic Network Programming – and Its Comparison with Genetic Programming. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, page 179, 2001. Poster online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2001/d02.pdf> [accessed 2008-06-15]. See also [1040, 1042].
 - [1042] Hironobu Katagiri, Kotaro Hirasawa, Jinglu Hu, Junichi Murata, and M. Kosaka. Network Structure Oriented Evolutionary Model: Genetic Network Programming – Its Comparison with Genetic Programming –. *Transactions of the Society of Instrument and Control Engineers*, 38(5):485–494, May 2002. ISSN 0453-4654. Accession number: 02A0504907. Journal Code: S0104A. See also [1040, 1041].
 - [1043] Kotaro Hirasawa, M. Okubo, Jinglu Hu, and Junichi Murata. Comparison between Genetic Network Programming (GNP) and Genetic Programming (GP). In *Proceedings of the 2001 Congress on Evolutionary Computation CEC2001*, pages 1276–1282, 2001. In proceedings [500].
 - [1044] Tadahiko Murata and Takashi Nakamura from the Department of Informatics, Kansai University, Takatsuki, 569-1095, Osaka, Japan. Multi-agent Cooperation Using Genetic Network Programming with Automatically Defined Groups. In *Genetic and Evolutionary Computation – GECCO 2004*, pages 712–714, 2004. In proceedings [559].
 - [1045] Tadahiko Murata and Takashi Nakamura. Developing Cooperation of Multiple Agents Using Genetic Network Programming with Automatically Defined Groups. In *Late Breaking Papers at the 2004 Genetic and Evolutionary Computation Conference*, 2004. In proceedings [560].
 - [1046] Shingo Mabu, Kotaro Hirasawa, and Jinglu Hu. Genetic Network Programming with Reinforcement Learning and its Performance Evaluation. In *Late Breaking Papers at the 2004 Genetic and Evolutionary Computation Conference*, 2004. In proceedings [560]. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2004/LBP036.pdf> [accessed 2008-06-16].
 - [1047] Julian Francis Miller and Peter Thomson from the Department of Computer Studies, Napier University, 219 Colinton Road, Edinburgh, EH14 1DJ, UK. Aspects of Digital Evolution: Geometry and Learning. In *Evolvable Sys-*

- tems: From Biology to Hardware*, volume 1478/1998 of *Lecture Notes in Computer Science (LNCS)*, pages 25–35. Springer Berlin/Heidelberg, 1998. ISBN 978-3-540-64954-0. doi:10.1007/BFb0057601. Online available at <http://citeseer.ist.psu.edu/40293.html> and <http://www.elec.york.ac.uk/intsys/users/jfm7/ices1998-miller-thomson.pdf> [accessed 2007-11-03].
- [1048] Julian Francis Miller. An empirical study of the efficiency of learning boolean functions using a Cartesian Genetic Programming approach. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1135–1142, 1999. In proceedings [572]. Online available at <http://citeseer.ist.psu.edu/miller99empirical.html> and <http://www.elec.york.ac.uk/intsys/users/jfm7/gecco1999b-miller.pdf> [accessed 2007-11-02].
- [1049] Julian Francis Miller and Peter Thomson. Cartesian Genetic Programming. In *Genetic Programming, Proceedings of EuroGP'2000*, pages 121–132, 2000. In proceedings [909]. Online available at <http://citeseer.ist.psu.edu/424028.html> and <http://www.elec.york.ac.uk/intsys/users/jfm7/cgp-eurogp2000.pdf> [accessed 2007-11-02].
- [1050] Janet Clegg, James Alfred Walker, and Julian Frances Miller. A New Crossover Technique for Cartesian Genetic Programming. In *GECCO'07: Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*, pages 1580–1587, 2007. In proceedings [552]. Online available at <http://doi.acm.org/10.1145/1276958.1277276> and <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2007/docs/p1580.pdf> [accessed 2007-11-01].
- [1051] Gwoing Tina Yu and Julian Francis Miller. Neutrality and the Evolvability of Boolean Function Landscape. In *EuroGP'01: Proceedings of the 4th European Conference on Genetic Programming*, pages 204–217, 2001. In proceedings [908]. Online available at <http://citeseer.ist.psu.edu/yu01neutrality.html> and http://www.cs.mun.ca/~tinayu/index_files/addr/public_html/neutrality.pdf [accessed 2007-11-03], see also http://www.cs.mun.ca/~tinayu/index_files/addr/public_html/neutralityTalk.pdf [accessed 2007-11-03].
- [1052] James Alfred Walker and Julian Francis Miller. Evolution and Acquisition of Modules in Cartesian Genetic Programming. In *7th European Conference on Genetic Programming, EuroGP 2004*, pages 187–197, 2004. In proceedings [905]. Online available at <http://www.cartesiangp.co.uk/papers/2004/wmEuroGP2004.pdf> and <http://www.elec.york.ac.uk/intsys/users/jfm7/eurogp2004.pdf> [accessed 2007-11-03].
- [1053] Peter John Angeline and Jordan Pollack. Evolutionary Module Acquisition. In *Proceedings of the Second Annual Conference on Evolutionary Programming*, pages 154–163, 1993. In proceedings [1191]. Online available at <http://www.demon.cs.brandeis.edu/papers/ep93.pdf> and <http://www.natural-selection.com/Library/1993/ep93.ps.Z> [accessed 2007-11-03].
- [1054] James Alfred Walker and Julian Francis Miller. Improving the Evolvability of Digital Multipliers Using Embedded Cartesian Genetic Programming and Product Reduction. In Juan Manuel Moreno, Jordi Madrenas, and Jordi Cosp, editors, *Evolvable Systems: From Biology to Hardware, Proceedings of 6th International Conference in Evolvable Systems (ICES 2005)*, volume 3637 of *Lecture Notes in Computer Science (LNCS)*, pages 131–142. Springer, September 12-14, 2005, Sitges, Spain. ISBN 3-540-28736-1. Online available at <http://www.cartesiangp.co.uk/papers/2005/wmices2005.pdf> [accessed 2007-11-03].
- [1055] James Alfred Walker and Julian Francis Miller. Investigating the Performance of Module Acquisition in Cartesian Genetic Programming. In *GECCO'05: Proceedings of the 2005 Conference on Genetic and Evolutionary Computation*, pages 1649–1656, 2005. doi:10.1145/1068009.1068287. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068287> and <http://www.cartesiangp.co.uk/papers/2005/wmgecco2005.pdf> [accessed 2007-11-03].

- [1056] Julian Francis Miller and Peter Thomson. Aspects of Digital Evolution: Evolvability and Architecture. In *Parallel Problem Solving from Nature – PPSN V*, pages 927–936, 1998. In proceedings [588]. Online available at <http://citeseer.ist.psu.edu/miller98aspects.html> and <http://www.elec.york.ac.uk/intsys/users/jfm7/ppsn1998-miller-thomson.pdf> [accessed 2007-11-03].
- [1057] Julian Francis Miller and Peter Thomson. Evolving Digital Electronic Circuits for Real-Valued Function Generation using a Genetic Algorithm. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 863–868, 1998. In proceedings [913]. Online available at <http://citeseer.ist.psu.edu/miller98evolving.html> [accessed 2007-11-03].
- [1058] Tatiana Kalganova and Julian Francis Miller. Evolving More Efficient Digital Circuits by Allowing Circuit Layout Evolution and Multi-Objective Fitness. In Adrian Stoica, Jason Lohn, and Didier Keymeulen, editors, *EH'99: The First NASA/DoD Workshop on Evolvable Hardware*, pages 54–63, July 19-21, 1999. IEEE Computer Society, Pasadena, California. ISBN 0-7695-0256-3. Online available at <http://citeseer.ist.psu.edu/kalganova99evolving.html> [accessed 2007-11-03].
- [1059] Simon Harding L. and Julian Francis Miller from the Department Of Electronics University of York, YO10 5DD, UK. Evolution of Robot Controller Using Cartesian Genetic Programming. In *Genetic Programming, Proceedings of EuroGP 2005*, pages 62–73, 2005. doi:10.1007/b107383. In proceedings [904]. Online available at <http://www.cartesiangp.co.uk/papers/2005/hmeurogp2005.pdf> [accessed 2008-08-12] 2007-11-03.
- [1060] James Alfred Walker and Julian Francis Miller. Embedded Cartesian Genetic Programming and the Lawnmower and Hierarchical-if-and-only-if Problems. In *GECCO'06: Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pages 911–918, 2006. In proceedings [553]. Online available at <http://www.cartesiangp.co.uk/papers/2006/wmgecco2006.pdf> [accessed 2007-11-03].
- [1061] James Alfred Walker and Julian Francis Miller from the Intelligent Systems Group, Department of Electronics, University of York, Heslington, York, YO10 5DD, UK. Predicting Prime Numbers Using Cartesian Genetic Programming. In *Proceedings of the 10th European Conference on Genetic Programming*, pages 205–216, 2007. doi:10.1007/978-3-540-71605-1_19. In proceedings [902].
- [1062] Bart Naudts. *Het meten van GA-hardheit – Measuring GA-hardness*. PhD thesis, Universitaire Instelling Antwerpen, Antwerpen, Netherlands, June 1998. Supervisor: Alain Verschoren. Online available at <http://citeseer.ist.psu.edu/185231.html> and <http://libra.msra.cn/paperdetail.aspx?id=322506> [accessed 2007-07-29].
- [1063] Walter Fontana. Algorithmic chemistry. In *Artificial Life II*, pages 159–210, 1990. In proceedings [1924].
- [1064] Christian W. G. Lasarczyk and Wolfgang Banzhaf. An Algorithmic Chemistry for Genetic Programming. In *Proceedings of the 8th European Conference on Genetic Programming*, pages 1–12, 2005. In proceedings [904]. Online available at <http://ls11-www.cs.uni-dortmund.de/downloads/papers/LaBa05.pdf> and <http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/eurogp-LasarczykB05.html> [accessed 2008-04-30].
- [1065] Christian W. G. Lasarczyk and Wolfgang Banzhaf. Total synthesis of algorithmic chemistries. In *GECCO 2005: Proceedings of the 2005 conference on Genetic and evolutionary computation*, volume 2, pages 1635–1640, 2005. In proceedings [555]. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2005/docs/p1635.pdf> and <http://doi.acm.org/10.1145/1068009.1068285> [accessed 2008-04-02].
- [1066] Wolfgang Banzhaf and Christian W. G. Lasarczyk. Genetic Programming of an Algorithmic Chemistry. In *Genetic Programming Theory and Practice II*, pages 175–190, 2005. doi:10.1007/0-387-23254-0-11. In proceedings [922]. Online available

- at <http://www.cs.mun.ca/~banzhaf/papers/algochem.pdf> and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/banzhaf_2004_GTP.html [accessed 2008-04-30].
- [1067] John Henry Holland and Judith S. Reitman. Cognitive Systems based on adaptive algorithms. In D. A. Waterman and F. Hayes-Roth, editors, *Pattern directed inference systems*, pages 313–329. Academic Press, New York, NY, 1978. Reprinted in [594]. See also [1196].
- [1068] John Henry Holland and Arthur W. Burks. *Adaptive computing system capable of learning and discovery*. Number 619349 filed on 1984-06-11 in United States Patent. United States Patent and Trademark Office, 1987. US Patent Issued on September 29, 1987, Current US Class 706/13, Genetic algorithm and genetic programming system 382/155, LEARNING SYSTEMS 706/62 MISCELLANEOUS, Foreign Patent References 8501601 WO Apr., 1985. Online available at <http://www.patentstorm.us/patents/4697242.html> and <http://www.freepatentsonline.com/4697242.html> [accessed 2008-04-01].
- [1069] William M. Spears and Kenneth Alan De Jong. Using genetic algorithms for supervised concept learning. In *Proceedings of the 2nd International IEEE Conference on Tools for Artificial Intelligence*, pages 335–341, 6-9 1990. IEEE Computer Society Press, Los Alamitos, CA, Herndon, VA. IEEE Cat. No. 90CH2915-7.
- [1070] Stephen Frederick Smith. *A learning system based on genetic adaptive algorithms*. PhD thesis, University of Pittsburgh, 1980. AAI8112638.
- [1071] Thomas Weise, Michael Zapf, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Rule-based Genetic Programming. In *Proceedings of BIONETICS 2007, 2nd International Conference on Bio-Inspired Models of Network, Information, and Computing Systems*, 2007. In proceedings [1936]. Online available at <http://www.it-weise.de/documents/files/WZG2007RBGP.pdf> [accessed 2008-08-12].
- [1072] Edsger Wybe Dijkstra from the Burroughs Corporation, Nuenen, The Netherlands. Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM*, 18(8):453–458, August 1975. ISSN 0001-0782. doi:10.1145/360933.360975. Online available at <http://www.cs.utexas.edu/users/EWD/ewd04xx/EWD472.PDF> and <http://doi.acm.org/10.1145/360933.360975> [accessed 2008-07-23].
- [1073] Will N. Browne and Charalambos Ioannides. Investigating scaling of an abstracted LCS utilising ternary and s-expression alphabets. In *GECCO'07: Proceedings of the 2007 GECCO conference companion on Genetic and evolutionary computation*, pages 2759–2764, 2007. doi:10.1145/1274000.1274067. In proceedings [551]. Online available at <http://portal.acm.org/citation.cfm?id=1274000.1274067> [accessed 2007-08-01].
- [1074] Pier Luca Lanzi and Alessandro Perrucci. Extending the Representation of Classifier Conditions Part II: From Messy Coding to S-Expressions. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 99)*, pages 345–352, 1999. In proceedings [572]. Online available at <http://webpace.elet.polimi.it/lanzi/papers//lanzi1999GECCOgp.pdf> [accessed 2007-08-01].
- [1075] Alan Mathison Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265, 1936. Online available at http://www.thocp.net/biographies/turing_alan.html and <http://www.abelard.org/turpap2/tp2-ie.asp> [accessed 2007-08-11].
- [1076] Astro Teller. Turing Completeness in the Language of Genetic Programming with Indexed Memory. In *Proceedings of the 1994 IEEE World Congress on Computational Intelligence*, volume 1, pages 136–141, June 27–29, 1994. IEEE Press, Orlando, Florida, USA. Online available at <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/Turing.ps> and <http://www.astroteller.net/work/pdfs/Turing.pdf> [accessed 2007-09-17].

- [1077] Nicholas Freitag McPhee and Ricardo Poli. Memory with Memory: Soft Assignment in Genetic Programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1235–1242, 2008. In proceedings [550].
- [1078] Mark A. Bedau from the Department of Philosophy, Reed College, 3023 SE Woodstock Blvd., Portland OR 97202, USA. Artificial life: organization, adaptation and complexity from the bottom up. *TRENDS in Cognitive Sciences*, 7(11):505–512, November 2003. Online available at <http://people.reed.edu/~mab/publications/papers/BedauTICS03.pdf> and <http://dx.doi.org/10.1016/j.tics.2003.09.012> [accessed 2007-12-13].
- [1079] Peter Dittrich, Jens Ziegler, and Wolfgang Banzhaf from the University of Dortmund, Department of Computer Science, Systems Analysis, D-44221 Dortmund, Germany, <http://ls11-www.cs.uni-dortmund.de>. Artificial Chemistries – A Review. *Artificial Life*, 7(3):225–275, 2001. Online available at <http://citeseer.ist.psu.edu/509901.html> and http://www.cs.mun.ca/~banzhaf/papers/alchemistry_review_MIT.pdf [accessed 2008-05-01].
- [1080] Tim J. Hutton from the Biomedical Informatics Unit, Eastman Dental Institute for Oral Health Care Sciences, University College London, 256 Gray’s Inn Road, London WC1X 8LD, UK. Evolvable Self-Replicating Molecules in an Artificial Chemistry. *Artificial Life*, 8(4):341–356, Winter 2002. ISSN 1064-5462. Online available at <http://www.sq3.org.uk/papers/evselfreps.pdf> [accessed 2008-05-01].
- [1081] Tim J. Hutton from the Biomedical Informatics Unit, Eastman Dental Institute for Oral Health Care Sciences, University College London, 256 Gray’s Inn Road, London WC1X 8LD, UK. Simulating Evolutions First Steps. In *Advances in Artificial Life, Proceedings of the 7th European Conference*, 2003. In proceedings [1937]. Online available at <http://www.sq3.org.uk/papers/ecal03.pdf> [accessed 2008-05-01].
- [1082] Tim J. Hutton from the Biomedical Informatics Unit, Eastman Dental Institute for Oral Health Care Sciences, University College London, 256 Gray’s Inn Road, London WC1X 8LD, UK. Evolvable Self-Reproducing Cells in a Two-Dimensional Artificial Chemistry. *Artificial Life*, 13(1):11–30, Spring 2007. ISSN 1064-5462. Online available at <http://www.sq3.org.uk/papers/cells2007.pdf> [accessed 2008-05-01].
- [1083] Lee Spector and Kilian Stoffel. Ontogenetic Programming. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 394–399, 1996. In proceedings [917]. Online available at <http://hampshire.edu/lspector/pubs/onto-gp96-e.ps> [accessed <http://citeseer.comp.nus.edu.sg/2137.html>] 2008-06-20.
- [1084] Kilian Stoffel and Lee Spector. High-Performance, Parallel, Stack-Based Genetic Programming. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 224–229, 1996. In proceedings [915]. Online available at <http://cognet.mit.edu/library/books/view?isbn=0262611279> and http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/stoffel_1996_hpsbGP.html [accessed 2008-06-20].
- [1085] Lee Spector. Autoconstructive Evolution: Push, PushGP, and Pushpop. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*, pages 137–146, 2001. In proceedings [568]. Online available at <http://citeseer.ist.psu.edu/445431.html> and <http://hampshire.edu/lspector/pubs/ace.pdf> [accessed 2007-12-24].
- [1086] Lee Spector and Alan Robinson from the Cognitive Science, Hampshire College, Amherst, MA 01002. Genetic Programming and Autoconstructive Evolution with the Push Programming Language. *Genetic Programming and Evolvable Machines*, 3(1):7–40, 2002. Received August 15, 2001; Revised November 26, 2001. Online available at <http://citeseer.ist.psu.edu/619906.html> and <http://hampshire.edu/lspector/pubs/push-gpem-final.pdf> [accessed 2007-12-25].
- [1087] Lee Spector, Chris Perry, Jon Klein, and Maarten Keijzer. Push 3.0 Programming Language Description. Technical Report HC-CSTR-2004-02, School of Cognitive Science, Hampshire College, Amherst, Massachusetts 01002, USA, September 9,

2004. Successor to the Push 2.0 Programming Language Description (<http://hampshire.edu/lspector/push2-description.html> [accessed 2007-12-25]), from which it borrows large chunks of text. Online available at http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/Spector_push3tr.html and http://www.hampshire.edu/cms_PDF/HC-CSTR-2004-02.pdf [accessed 2007-12-25].
- [1088] Lee Spector, Jon Klein, and Maarten Keijzer. The Push3 execution stack and the evolution of control. In *GECCO 2005: Proceedings of the 2005 conference on Genetic and evolutionary computation*, volume 2, pages 1689–1696, 2005. In proceedings [555]. Online available at <http://doi.acm.org/10.1145/1068009.1068292> and <http://hampshire.edu/lspector/pubs/push3-gecco2005.pdf> [accessed 2007-12-25].
- [1089] Raphael Crawford-Marks and Lee Spector. Size Control Via Size Fair Genetic Operators In The PushGP Genetic Programming System. In *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 733–739, 2002. In proceedings [563]. Online available at <http://citeseer.ist.psu.edu/608051.html> and <http://alum.hampshire.edu/~rpc01/gp234.pdf> [accessed 2007-12-25].
- [1090] Alan Robinson. Genetic Programming: Theory, Implementation, and the Evolution of Unconstrained Solutions. Master’s thesis, Cognitive Science, Hampshire College, Amherst, MA 01002, May 2003. Division III Thesis. Committee Lee Spector, Jaime Davila, Mark Feinstein. Online available at <http://hampshire.edu/lspector/robinson-div3.pdf> and <http://citeseer.ist.psu.edu/498673.html> [accessed 2007-12-25].
- [1091] Lee Spector. Adaptive populations of endogenously diversifying Pushpop organisms are reliably diverse. In Standish, Abbass, and Bedau, editors, *ICAL 2003: Proceedings of the eighth international conference on Artificial life*, pages 142–145. MIT Press, Cambridge, MA, USA, December 9–12, 2002, University of New South Wales, Sydney, NSW, Australia. ISBN 0-262-69281-3. Online available at <http://www.alife.org/alife8/proceedings/sub962.pdf> and <http://citeseer.ist.psu.edu/627720.html> [accessed 2007-12-25].
- [1092] Christian F. Tschudin from the Computer Science Department, University of Basel, Bernoullistrasse 16, CH-4056 Basel, Switzerland. Fraglets – a Metabolic Execution Model for Communication Protocols. In *Proceedings of 2nd Annual Symposium on Autonomous Intelligent Networks and Systems (AINS)*, June 30–July 3, 2003, SRI International, Menlo Park, CA, USA. Online available at <http://cn.cs.unibas.ch/pub/doc/2003-ains.pdf> and <http://path.berkeley.edu/ains/final/007%20-%2022-tschudin.pdf> [accessed 2008-05-02]. See also <http://www.fraglets.net/> [accessed 2007-09-17].
- [1093] Emil Leon Post. Formal reductions of the general combinatorial decision problem. *American Journal of Mathematics*, 65(2):197–215, 1943, Baltimore, MD, USA. ISSN 0002-9327, 1080-6377. doi:10.2307/2371809.
- [1094] Stephen Cole Kleene. On notation for ordinal numbers. *The Journal of Symbolic Logic*, 3(4):150–155, 1938.
- [1095] Lidia A. R. Yamamoto, Daniel Schreckling, and Thomas Meyer. Self-Replicating and Self-Modifying Programs in Fraglets. In *Proceedings of BIONETICS 2007, 2nd International Conference on Bio-Inspired Models of Network, Information, and Computing Systems*, 2007. In proceedings [1936]. Online available at <http://cn.cs.unibas.ch/people/ly/doc/bionetics2007-ysm.pdf> [accessed 2008-05-04].
- [1096] Martin Davis. *Engines of Logic*. W.W. Norton & Company, London, 2000. ISBN 039-3322-297.
- [1097] Michael Sipser. *Introduction to the Theory of Computation*. PWS Publishing, second edition, 2006. ISBN 053-4947-28X.
- [1098] David Hilbert and Wilhelm Friedrich Ackermann. *Grundzüge der theoretischen Logik*, volume XXVII of *Die Grundlehren der mathematischen Wissenschaften in Einzel-*

- darstellungen*. Springer, Berlin, August 1928. ISBN ISBN 0-8218-2024-9. English version in [1099].
- [1099] David Hilbert and Wilhelm Friedrich Ackermann. *Principles of Mathematical Logic*. AMS Chelsea Publishing, 1950. ISBN 0-8218-2024-9, 978-0-8218-2024-7. Edited by Robert E. Luce, translated by Lewis M. Hammond, George G. Leckie, and F. Steinhardt. Partly online available at <http://books.google.com/books?id=7E34DeXok9gC> [accessed 2008-03-25]. Original German version: [1098].
 - [1100] Alonzo Church. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2):345–363, April 1936.
 - [1101] Alonzo Church. A note on the Entscheidungsproblem. *Journal of Symbolic Logic*, 1(1):40–41, March 1936.
 - [1102] Alan Mathison Turing. On computable numbers, with an application to the Entscheidungsproblem, Errata. *Proceedings of the London Mathematical Society*, 43:544–546, 1937. Online available at the same location as [1075].
 - [1103] Astro Teller from the Carnegie Mellon University, Department of Computer Science, 4019 Winterburn Ave, Pittsburgh, PA 15207, USA. Genetic Programming, Indexed memory, the Halting problem, and other curiosities. In *Proceedings of the 7th annual Florida Artificial Intelligence Research Symposium*, pages 270–274, May 1994. IEEE Press, Pensacola, Florida, USA. Online available at <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/Curiosities.ps> and <http://citeseer.ist.psu.edu/256509.html> [accessed 2008-06-16].
 - [1104] Boris Beizer. *Software Testing Techniques*. International Thomson Computer Press, second edition, June 1990. ISBN 185-0328-803, 978-1850328803.
 - [1105] Mark Fewster and Dorothy Graham. *Software Test Automation*. Addison-Wesley Professional, August 25, 1999. ISBN 020-1331-403.
 - [1106] D. Gelperin and B. Hetzel. The Growth of Software Testing. *Communications of the ACM*, 31:687–695, 1988. ISSN 0001-0782. Online available at <http://doi.acm.org/10.1145/62959.62965> [accessed 2007-09-15].
 - [1107] Cem Kaner, Jack Falk, and Hung Quoc Nguyen. *Testing Computer Software*. John Wiley and Sons, second edition, April 12, 1993. ISBN 047-1358-460, 978-0471358466.
 - [1108] David Jackson. Parsing and translation of expressions by genetic programming. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1681–1688, 2005. doi:10.1145/1068009.1068291. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068291> [accessed 2008-05-04].
 - [1109] Sean Luke. Code Growth is Not Caused by Introns. In *Late Breaking Papers at the 2000 Genetic and Evolutionary Computation Conference*, pages 228–235, 2000. In proceedings [571]. Online available at <http://citeseer.ist.psu.edu/300709.html> and <http://www.cs.gmu.edu/~sean/papers/intronpaper.pdf> [accessed 2007-09-07].
 - [1110] Tobias Blickle and Lothar Thiele. Genetic Programming and Redundancy. In J. Hopf, editor, *Genetic Algorithms within the Framework of Evolutionary Computation (Workshop at KI-94, Saarbrücken)*, pages 33–38, 1994. Max-Planck-Institut für Informatik (MPI-I-94-241), Im Stadtwald, Building 44, D-66123 Saarbrücken, Germany. Online available at <http://citeseer.ist.psu.edu/44816.html> and <http://www.tik.ee.ethz.ch/~tec/publications/bt94/GPandRedundancy.ps.gz> [accessed 2007-09-07].
 - [1111] William B. Langdon and Riccardo Poli. Fitness Causes Bloat: Mutation. In *Proceedings of the First European Workshop on Genetic Programming*, pages 37–48, 1998. In proceedings [911]. Online available at <http://citeseer.ist.psu.edu/langdon98fitness.html> and ftp://ftp.cwi.nl/pub/W.B.Langdon/papers/WBL.euro98_bloatm.ps.gz [accessed 2007-09-09].
 - [1112] Walter Alden Tackett. *Recombination, Selection, and the Genetic Construction of Computer Programs*. PhD thesis, Faculty of the Graduate School of the University of

- Southern California, Los Angeles, CA, USA, April 1994. Online available at <http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/ftp.io.com/papers/watphd.tar.Z> [accessed 2007-09-07].
- [1113] Peter W. H. Smith and Kim Harries from the Computer Science Department, City University, Northampton Square, London EC1V 0HB, UK. Code growth, explicitly defined introns, and alternative selection schemes. *Evolutionary Computation*, 6(4): 339–360, Winter 1998. ISSN 1063-6560, 1530-9304. doi:10.1162/evco.1998.6.4.339. Online available at <http://citeseer.ist.psu.edu/harries98code.html> and [accessed 2008-07-22].
- [1114] Nicholas Freitag McPhee and Justin Darwin Miller. Accurate Replication in Genetic Programming. In *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 303–309, 1995. In proceedings [747]. Online available at http://www.mrs.umn.edu/~mcphee/Research/Accurate_replication.ps and <http://citeseer.ist.psu.edu/mcphee95accurate.html> [accessed 2007-09-07].
- [1115] Justinian Rosca. Generality Versus Size in Genetic Programming. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 381–387, 1996. In proceedings [917]. Online available at <http://citeseer.ist.psu.edu/75165.html> and <ftp://ftp.cs.rochester.edu/pub/u/rosca/gp/96.gp.ps.gz> [accessed 2007-09-07].
- [1116] Tobias Blickle. *Theory of Evolutionary Algorithms and Application to System Synthesis*. PhD thesis, ETH Zurich, Verlag der Fachvereine Hochschulverlag AG of ETH Zurich, November 1996. ETH-Thesis number 11894. Examiner: Lothar Thiele, Hans-Paul Schwefel. Online available at <http://www.tik.ee.ethz.ch/~tec/publications/bli97a/diss.ps.gz> and <http://www.handshake.de/user/blickle/publications/diss.pdf> [accessed 2007-09-07].
- [1117] William B. Langdon, Terry Soule, Riccardo Poli, and James A. Foster. The Evolution of Size and Shape. In *Advances in Genetic Programming 3*, chapter 8, pages 163–190. The MIT Press, 1999. In collection [930]. Online available at <http://www.cs.bham.ac.uk/~wbl/aigp3/ch08.ps.gz> and <http://citeseer.ist.psu.edu/langdon99evolution.html> [accessed 2007-09-07].
- [1118] Peter Nordin, Frank Francone, and Wolfgang Banzhaf. Explicitly defined introns and destructive crossover in genetic programming. In Peter John Angeline and Jr Kenneth E. Kinneer, editors, *Advances in genetic programming 2*, pages 111–134. The MIT Press, 1996. In collection [929] and also [1030]. Online available at <ftp://lumpi.informatik.uni-dortmund.de/pub/biocomp/papers/ML95.ps.gz> [accessed 2007-08-17].
- [1119] Terence Soule and James A. Foster. Removal Bias: a New Cause of Code Growth in Tree Based Evolutionary Programming. In *1998 IEEE International Conference on Evolutionary Computation*, pages 781–186, 1998. In proceedings [503]. Online available at <http://citeseer.ist.psu.edu/313655.html> [accessed 2007-09-07].
- [1120] Stefan Bleuler, Martin Brack, Lothar Thiele, and Eckart Zitzler. Multiobjective Genetic Programming: Reducing Bloat Using SPEA2. In *Proceedings of the 2001 Congress on Evolutionary Computation CEC2001*, pages 536–543, 2001. In proceedings [500]. Online available at <http://citeseer.ist.psu.edu/443099.html> and <ftp://ftp.tik.ee.ethz.ch/pub/people/zitzler/BBTZ2001b.ps.gz> [accessed 2007-09-07].
- [1121] Anikó Ekárt and S. Z. Németh from the Computer and Automation Research Institute, Hungarian Academy of Sciences, 1518 Budapest, P.O.B. 63, Hungary. Selection Based on the Pareto Nondomination Criterion for Controlling Code Growth in Genetic Programming. *Genetic Programming and Evolvable Machines*, 2(1):61–73, March 2001. ISSN 1389-2576 (Print) 1573-7632 (Online). doi:10.1023/A:1010070616149. Online available at <http://www.springerlink.com/content/j70v0303523v6x65/fulltext.pdf> [accessed 2008-05-04].
- [1122] Sean Luke and Liviu Panait. A comparison of bloat control methods for genetic programming. *Evolutionary Computation*, 14(3):309–344, 2006. ISSN 1063-6560.

- [1123] Riccardo Poli. A Simple but Theoretically-Motivated Method to Control Bloat in Genetic Programming. In *Proceedings of 6th European Conference on Genetic Programming, EuroGP 2003*, pages 204–217, 2003. In proceedings [906].
- [1124] Ingo Rechenberg. *Cybernetic Solution Path of an Experimental Problem*. Royal Aircraft Establishment, August 1965. Library Translation 1122, Farnborough. Reprinted in [594].
- [1125] Thomas Bäck, Frank Hoffmeister, and Hans-Paul Schwefel. A survey of evolution strategies. In *Proceedings of the 4th International Conference on Genetic Algorithms ICGA'91*, pages 2–9, 1991. In proceedings [749]. Online available at <http://de.scientificcommons.org/50038> and <http://citeseer.ist.psu.edu/19412.html> [accessed 2007-08-27].
- [1126] Hans-Georg Beyer and Hans-Paul Schwefel from the Department of Computer Science XI, University of Dortmund, Joseph-von-Fraunhoferstrasse 20, D-44221 Dortmund, Germany. Evolution strategies – A comprehensive introduction. *Natural Computing: an international journal*, 1(1):3–52, March 2002. ISSN 1567-7818 (Print) 1572-9796 (Online). doi:10.1023/A:1015059928466. Online available at <http://dx.doi.org/10.1023/A:1015059928466> and <http://www.springerlink.com/content/2311qapbrwgrcyey/> [accessed 2007-08-27].
- [1127] Hans-Georg Beyer. *The theory of evolution strategies*. Natural Computing Series. Springer-Verlag New York, Inc., New York, NY, USA, March 27, 2001. ISBN 978-3-540-67297-5, 3540672974. Partly online available at <http://books.google.de/books?id=8tbInLufkTMC> [accessed 2008-06-07]. Series editors: G. Rozenberg, Thomas Bäck, Agoston E. Eiben, J.N. Kok, and H.P. Spalink.
- [1128] Nikolaus Hansen, Andreas Ostermeier, and Andreas Gawelczyk. On the Adaptation of Arbitrary Normal Mutation Distributions in Evolution Strategies: The Generating Set Adaptation. In *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 57–64, 1995. In proceedings [747]. Online available at <http://citeseer.ist.psu.edu/261932.html> [accessed 2007-08-27].
- [1129] Silja Meyer-Nieberg and Hans-Georg Beyer. Self-Adaptation in Evolutionary Algorithms. In *Parameter Setting in Evolutionary Algorithms*. Springer, 2007. In collection [609]. Online available at <http://citeseer.ist.psu.edu/741946.html> [accessed 2007-07-28].
- [1130] Nikolaus Hansen and Andreas Ostermeier. Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In *Proceedings of IEEE International Conference on Evolutionary Computation*, pages 312–317, 1996. In proceedings [505]. Online available at <http://citeseer.ist.psu.edu/hansen96adapting.html> and <http://www.bionik.tu-berlin.de/ftp-papers/CMAES.ps.Z> [accessed 2007-09-20].
- [1131] Nikolaus Hansen and Andreas Ostermeier from the Fachgebiet für Bionik und Evolutionstechnik, Technische Universität Berlin. Convergence Properties of Evolution Strategies with the Derandomized Covariance Matrix Adaption: The $(\mu/\mu_I, \lambda)$ -CMA-ES. In Hans-Jürgen Zimmermann, editor, *Eufit'97 – 5th European Congress on Intelligent Techniques and Soft Computing*, pages 650–654. Verlag Mainz, Wissenschaftsverlag, 1997, Aachen. Online available at <http://citeseer.ist.psu.edu/356709.html> [accessed 2007-08-27].
- [1132] Nikolaus Hansen and Andreas Ostermeier from the Fachgebiet für Bionik und Evolutionstechnik, Technische Universität Berlin. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation*, 9(2):159–195, 2001. Online available at <http://www.bionik.tu-berlin.de/user/niko/cmaartic.pdf> and <http://citeseer.ist.psu.edu/hansen01completely.html> [accessed 2007-08-27].
- [1133] G. A. Jastrebski and D. V. Arnold. Improving evolution strategies through active covariance matrix adaptation. In *Proceedings of IEEE World Congress on Computa-*

- tional Intelligence*, pages 9719–9726. IEEE Press, Piscataway, New Jersey, July 2006, Vancouver, BC.
- [1134] Hsien-Da Huang, Jih Tsung Yang, Shu Fong Shen, and Jorng-Tzong Horng. An Evolution Strategy to Solve Sports Scheduling Problems. In *Proceedings of the Genetic and Evolutionary Computation Conference GECCO-99*, volume 1, page 943, 1999. In proceedings [572].
 - [1135] Peter Roosen and Fred Meyer. Determination of Chemical Equilibria by means of an Evolution Strategy. In *PPSN-II, Parallel problem solving from nature 2*, pages 411–420, 1992. In proceedings [591].
 - [1136] Vincenzo Cutello, Giuseppe Narzisi, and Giuseppe Nicosia. A Class of Pareto Archived Evolution Strategy Algorithms Using Immune Inspired Operators for Ab-Initio Protein Structure Prediction. In *Applications on Evolutionary Computing, Proceedings of EvoWorkshops 2005*, pages 54–63, 2005. doi:10.1007/b106856. In proceedings [539].
 - [1137] Michael Emmerich, Monika Grötzner, Bernd Groß, and Martin Schütz. Mixed-Integer Evolution Strategy for Chemical Plant Optimization. In I.C. Parmee, editor, *Evolutionary Design and Manufacture – Selected papers from ACDM’00*, pages 55–67, April 2000, Plymouth, UK. Springer, London, New York. Online available at <http://citeseer.ist.psu.edu/359756.html> and <http://www.liacs.nl/~emmerich/pdf/EGG+00.pdf> [accessed 2007-08-27].
 - [1138] P. O’Brien, D. Corcoran, and D. Lowry. An evolution strategy to estimate emission source distributions on a regional scale from atmospheric observations. *Atmospheric Chemistry & Physics Discussions*, 3:1333–1366, February 2003. Provided by the Smithsonian/NASA Astrophysics Data System. Online available at <http://www.atmos-chem-phys-discuss.net/3/1333/2003/acpd-3-1333-2003.html> [accessed 2007-08-27].
 - [1139] Volker Nissen and Matthias Krause. Constrained Combinatorial Optimization with an Evolution Strategy. In *Proceedings of Fuzzy Logik, Theorie und Praxis, 4. Dortmunder Fuzzy-Tage*, pages 33–40, June 1994, Dortmund, Germany. Springer-Verlag, London, UK. ISBN 3-540-58649-0.
 - [1140] Hans-Georg Beyer. Some Aspects of the ‘Evolution Strategy’ for Solving TSP-Like Optimization Problems Appearing at the Design Studies on a $0.5\text{TeV}e^+e^-$ -Linear Collider. In *Parallel problem solving from nature 2*, pages 361–370, 1992. In proceedings [591].
 - [1141] Hans-Georg Beyer. Design optimization of a linear accelerator using evolution strategy: solving a TSP-like optimization problem. In *Handbook of Evolutionary Computation*, pages G4.2:1–8. Institute of Physics Publishing and Oxford University Press, 1997. In collection [410].
 - [1142] Robert E. Keller, Wolfgang Banzhaf, Jörn Mehnen, and Klaus Weinert. CAD surface reconstruction from digitized 3D point data with a genetic programming/evolution strategy hybrid. In *Advances in genetic programming: volume 3*, chapter 3, pages 41–65. The MIT Press, 1999. In collection [930]. Online available at <http://www.cs.bham.ac.uk/~wbl/aigp3/ch03.pdf> [accessed 2007-08-18].
 - [1143] Klaus Weinert and Jörn Mehnen. Discrete NURBS-Surface Approximation using an Evolutionary Strategy. Technical Report Sonderforschungsbereich (SFB) 531, Department of Machining Technology, University of Dortmund, Germany, October 2001. Online available at <http://citeseer.ist.psu.edu/566038.html> [accessed 2007-08-27].
 - [1144] Thomas Bäck and Martin Schütz. Evolution Strategies for Mixed-Integer Optimization of Optical Multilayer Systems. In *Evolutionary Programming IV: Proceedings of the Fourth Annual Conference on Evolutionary Programming*, pages 33–51, 1995. In proceedings [1189]. Online available at <http://citeseer.ist.psu.edu/101945.html> and <http://de.scientificcommons.org/114486> [accessed 2007-08-27].

- [1145] Rui Li, Michael T.M. Emmerich, Jeroen Eggermont, and Ernst G.P. Bovenkamp. Mixed-integer optimization of coronary vessel image analysis using evolution strategies. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1645–1652, 2006. doi:10.1145/1143997.1144268. In proceedings [553]. Online available at <http://doi.acm.org/10.1145/1143997.1144268> [accessed 2007-08-27].
- [1146] Hans-Paul Schwefel. *Numerical Optimization of Computer Models*. John Wiley and Sons Ltd, New York, NY, USA, June 17, 1981. ISBN 0471099880, 978-0471099888.
- [1147] Frank Hoffmeister and Hans-Paul Schwefel. KORR 2.1 – An implementation of a $(\mu + \lambda)$ -evolution strategy. Technical Report, Universität Dortmund, Fachbereich Informatik, November 1990. Interner Bericht.
- [1148] Yoshiyuki Matsumura, Kazuhiro Ohkura, and Kanji Ueda. Advantages of Global Discrete Recombination in $(\mu/\mu, \lambda)$ -Evolution Strategies. In *Proceedings of the 2002 Congress on Evolutionary Computation CEC2002*, volume 2, pages 1848–1853, 2002. doi:10.1109/CEC.2002.1004524. In proceedings [499].
- [1149] Hannes Geyer, Peter Ulbig, and Siegfried Schulz. Encapsulated Evolution Strategies for the Determination of Group Contribution Model Parameters in Order to Predict Thermodynamic Properties. In *PPSN V: Proceedings of the 5th International Conference on Parallel Problem Solving from Nature*, pages 978–987, 1998. In proceedings [588].
- [1150] Hannes Geyer, Peter Ulbig, and Siegfried Schulz. Use of evolutionary algorithms for the calculation of group contribution parameters in order to predict thermodynamic properties – Part 2: Encapsulated evolution strategies. *Computers and Chemical Engineering*, 23(7):955–973, July 1, 1999. doi:10.1016/S0098-1354(99)00270-7.
- [1151] Hannes Geyer, Peter Ulbig, and Siegfried Schulz. Verschachtelte Evolutionssstrategien zur Optimierung nichtlinearer verfahrenstechnischer Regressionsprobleme. *Chemie Ingenieur Technik*, 72(4):369–373, 2000. ISSN 1522-2640. doi:10.1002/1522-2640(200004)72:4<369::AID-CITE369>3.0.CO;2-W. Online available at <http://www3.interscience.wiley.com/cgi-bin/fulltext/76500452/PDFSTART> [accessed 2007-08-27].
- [1152] Kenneth V. Price, Rainer M. Storn, and Jouni A. Lampinen. *Differential Evolution – A Practical Approach to Global Optimization*. Natural Computing Series. Springer, New York, October 2005. ISBN 978-3-540-20950-8. Series editors: G. Rozenberg, Thomas Bäck, Agoston E. Eiben, J.N. Kok, and H.P. Späink. Partly online available at <http://books.google.de/books?id=S67vX-KqVqUC> [accessed 2008-07-23].
- [1153] Vitaliy Feoktistov. *Differential Evolution – In Search of Solutions*, volume 5 of *Springer Optimization and Its Applications*. Springer, December 2006. ISBN 978-0-387-36895-5.
- [1154] Efrñn Mezura-Montes, Jesús Velázquez-Reyes, and Carlos Artemio Coello Coello. A comparative study of differential evolution variants for global optimization. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 485–492, 2006. doi:10.1145/1143997.1144086. In proceedings [553]. Online available at <http://portal.acm.org/citation.cfm?id=1144086> and <http://delta.cs.cinvestav.mx/~ccoello/2006.html> [accessed 2007-08-13].
- [1155] Janez Brest, Viljem Žumer, and Mirjam Sepesy Maučec. Control Parameters in Self-Adaptive Differential Evolution. In *Proceedings of the Second International Conference on Bioinspired Optimization Methods and their Application, BIOMA 2006*, pages 35–44, 2006. In proceedings [491]. Online available at <http://labraj.uni-mb.si/index.php/Bibliografija> [accessed 2007-08-13].
- [1156] Jouni Lampinen and Ivan Zelinka. On stagnation of the differential evolution algorithm. In *Proceedings of MENDEL 2000, 6th International Mendel Conference on Soft Computing*, pages 76–83, 2000. In proceedings [379]. Online avail-

- able at <http://citeseer.ist.psu.edu/317991.html> and <http://www.lut.fi/~jlampine/MEND2000.ps> [accessed 2007-08-13].
- [1157] P. Besson, J.-M. Vesin, V. Popovici, and M. Kunt from the EPFL. Differential Evolution applied to a multimodal information theoretic optimization problem. In *Applications of Evolutionary Computing, Proceedings of the 8th European Workshop on Evolutionary Computation in Image Analysis and Signal Processing (evoIASP)*, pages 505–509, 2006. In proceedings [538].
- [1158] Rainer Storn and Kenneth Price. Differential Evolution – a simple and efficient adaptive scheme for global optimization over continuous spaces. Technical Report TR-95-012, International Computer Science Institute, 1947 Center Street, Berkeley, CA 94704, Berkeley, CA, 1995. Online available at <http://citeseer.ist.psu.edu/182432.html> and <http://http.icsi.berkeley.edu/~storn/TR-95-012.pdf> [accessed 2007-08-13].
- [1159] Yu Gao and Yong-Jun Wang. A Memetic Differential Evolutionary Algorithm for High Dimensional Functions’ Optimization. In *Third International Conference on Natural Computation (ICNC 2007), Part 4*, pages 188–192, 2007. doi:10.1109/ICNC.2007.60. In proceedings [313].
- [1160] Professor Kaelo and M. Montaz Ali from the School of Computational and Applied Mathematics, Witwatersrand University, Private-Bag-3, Wits-2050, Johannesburg, South Africa. A numerical study of some modified differential evolution algorithms. *European Journal of Operational Research*, 169(3):1176–1184, March 16, 2006. ISSN 0377-2217. doi:10.1016/j.ejor.2004.08.047. Online available at <http://witsetd.wits.ac.za:8080/dspace/bitstream/123456789/1771/2/PaperA.pdf> and <http://dx.doi.org/10.1016/j.ejor.2004.08.047> [accessed 2008-07-23].
- [1161] Professor Kaelo and M. Montaz Ali from the School of Computational and Applied Mathematics, Witwatersrand University, Wits-2050, Johannesburg, Republic of South Africa. Differential evolution algorithms using hybrid mutation. *Computational Optimization and Applications*, 37(2):231–246, June 2007. ISSN 0926-6003 (Print) 1573-2894 (Online). doi:10.1007/s10589-007-9014-3.
- [1162] Professor Kaelo. *Some population-set based methods for unconstrained global optimization*. PhD thesis, School of Computational and Applied Mathematics, Witwatersrand University, Private-Bag-3, Wits-2050, Johannesburg, South Africa, 2005. Online available at <http://hdl.handle.net/123456789/1771> [accessed 2008-07-23].
- [1163] M. Montaz Ali, Charoenchai Khompatraporn, and Zelda B. Zabinsky. A Numerical Evaluation of Several Stochastic Algorithms on Selected Continuous Global Optimization Test Problems. *Journal of Global Optimization*, 31(4):635–672, April 2005. ISSN 0925-5001 (Print) 1573-2916 (Online). doi:10.1007/s10898-004-9972-2. Online available at <http://www.springerlink.com/content/q825578116kn6k72/fulltext.pdf> [accessed 2008-07-23].
- [1164] Jouni Lampinen and Ivan Zelinka. Mechanical Engineering Design Optimization by Differential Evolution. In *New Ideas in Optimization*, pages 127–146. McGraw-Hill, 1999. In collection [400].
- [1165] Andreas Nearchou and Sotiris Omirou. Differential evolution for sequencing and scheduling optimization. *Journal of Heuristics*, 12(6):395–411(17), December 2006. Online available at <http://www.springerlink.com/content/p6761p070470k448/fulltext.pdf> [accessed 2007-08-13].
- [1166] Feng-Sheng Wang and Houng-Jhy Jang. Parameter Estimation of A Bioreaction Model by Hybrid Differential Evolution. In *Proceedings of the 2000 Congress on Evolutionary Computation*, pages 410–417, 2000. In proceedings [501].
- [1167] Colin C. Seaton and Maryjane Tremayne. Differential evolution: crystal structure determination of a triclinic polymorph of adipamide from powder diffraction data. *Chemical Communications (Camb.)*, 880(8):1, April 2002. On-

- line available at <http://www.rsc.org/Publishing/Journals/CC/article.asp?doi=b200436d> [accessed 2007-08-13].
- [1168] Maryjane Tremayne, Colin C. Seaton, and Christopher Glidewell. Structures of three substituted arenesulfonamides from X-ray powder diffraction data using the differential evolution technique. *Acta Crystallographica Section B: Structural Science*, 58: 823–834, October 2002. ISSN 0108-7681. Online available at <http://scripts.iucr.org/cgi-bin/paper?S0108768102011928> [accessed 2007-08-13].
 - [1169] Samantha Y. Chong and Maryjane Tremayne. Combined optimization using cultural and differential evolution: application to crystal structure solution from powder diffraction data. *Chemical Communication*, 39:4078–4080, October 2006. ISSN 1359-7345. Online available at <http://www.rsc.org/publishing/journals/CC/article.asp?doi=b609138e> [accessed 2007-08-13].
 - [1170] Yung-Chien Lin, Kao-Shing Hwang, and Feng-Sheng Wang. Plant Scheduling and Planning Using Mixed-Integer Hybrid Differential Evolution with Multiplier Updating. In *Proceedings of the 2000 Congress on Evolutionary Computation*, volume 1, pages 593–600, 2000. In proceedings [501].
 - [1171] Rainer Storn. On the Usage of Differential Evolution for Function Optimization. In M. Smith, M. Lee, J. Keller, and J. Yen, editors, *1996 Biennial Conference of the North American Fuzzy Information Processing Society*, pages 519–523, 1996. IEEE Press, Piscataway, NJ.
 - [1172] Rainer Storn. Differential Evolution Design of an IIR-Filter with Requirements for Magnitude and Group Delay. Technical Report TR-95-026, International Computer Science Institute, 1947 Center Street, Berkeley, CA 94704-1198, Berkeley, CA, 1995. Online available at <http://citeseer.ist.psu.edu/storn95differential.html> and <http://www.icsi.berkeley.edu/ftp/pub/techreports/1995/tr-95-026.pdf> [accessed 2007-08-13].
 - [1173] Rainer Storn. Designing Digital Filters with Differential Evolution. In *New Ideas in Optimization*, pages 109–125. McGraw-Hill Education, 1999. In collection [400].
 - [1174] David B. Fogel. *System Identification through Simulated Evolution: A Machine Learning Approach to Modeling*. Ginn Press, Needham Heights, MA, June 1991. ISBN 978-0536579430.
 - [1175] David B. Fogel. Evolving a checkers player without relying on human experience. *Intelligence*, 11(2):20–27, 2000. ISSN 1523-8822. Online available at <http://doi.acm.org/10.1145/337897.337996> [accessed 2007-08-27].
 - [1176] David B. Fogel. *Blondie24: playing at the edge of AI*. The Morgan Kaufmann Series in Artificial Intelligence. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, September 2001. ISBN 978-1558607835.
 - [1177] Lawrence J. Fogel and David B. Fogel. A Preliminary Investigation on Extending Evolutionary Programming to Include Self-Adaptation on Finite State. *Informatica (Slovenia)*, 18(4), 1994.
 - [1178] Vincent W. Porto, David B. Fogel, and Lawrence J. Fogel. Alternative Neural Network Training Methods. *IEEE Expert: Intelligent Systems and Their Applications*, 10(3):16–22, 1995. ISSN 0885-9000. doi:10.1109/64.393138.
 - [1179] Daniel K. Gehlhaar, Gennady Verkhivker, Paul A. Rejto, David B. Fogel, Lawrence J. Fogel, and Stephan T. Freer. Docking Conformationally Flexible Small Molecules into a Protein Binding Site through Evolutionary Programming. In *Evolutionary Programming IV: Proceedings of the Fourth Annual Conference on Evolutionary Programming*, 1995. In proceedings [1189].
 - [1180] Bruce S. Duncan and Arthur J. Olson. Applications of Evolutionary Programming for the Prediction of Protein-Protein Interactions. In *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming*, pages 411–417, 1996. In proceedings [1188].

- [1181] Daniel K. Gehlhaar and David B. Fogel. Tuning Evolutionary Programming for Conformationally Flexible Molecular Docking. In *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming*, pages 419–429, 1996. In proceedings [1188].
- [1182] Jong-Hwan Kim and Jeong-Yul Jeon. Evolutionary Programming-Based High-Precision Controller Design. In *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming*, pages 73–81, 1996. In proceedings [1188].
- [1183] Kevin M. Nelson. A Comparison of Evolutionary Programming and Genetic Algorithms for Electronic Part Placement. In *Evolutionary Programming IV: Proceedings of the Fourth Annual Conference on Evolutionary Programming*, pages 503–519, 1995. In proceedings [1189].
- [1184] Manish Sarkar. Evolutionary Programming-Based Fuzzy Clustering. In *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming*, pages 247–256, 1996. In proceedings [1188].
- [1185] Jong-Hwan Kim and Hyun-Sik Shim. Evolutionary Programming-Based Optimal Robust Locomotion Control of Autonomous Mobile Robots. In *Evolutionary Programming IV: Proceedings of the Fourth Annual Conference on Evolutionary Programming*, pages 631–644, 1995. In proceedings [1189].
- [1186] V. William Porto, N. Saravanan, D. Waagen, and Agoston E. Eiben, editors. *Proceedings of the 7th International Conference on Evolutionary Programming VII, EP98, in cooperation with IEEE Neural Networks Council*, volume 1447/1998 of *Lecture Notes in Computer Science (LNCS)*, March 25–27, 1998, Mission Valley Marriott, San Diego, California, USA. Springer Verlag, Heidelberg, Germany. ISBN 978-3540648918.
- [1187] Peter John Angeline, Robert G. Reynolds, John Robert McDonnell, and Russel C. Eberhart, editors. *Proceedings of the 6th International Conference on Evolutionary Programming*, volume 1213/1997 of *Lecture Notes in Computer Science (LNCS)*, 1997, Indianapolis, Indiana, USA. Springer Verlag Telos. ISBN 978-3540627883.
- [1188] Lawrence J. Fogel, Peter John Angeline, and Thomas Bäck, editors. *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming (EP'96)*, February 1996, San Diego, California, USA. MIT Press, Cambridge, Massachusetts. ISBN 0-262-06190-2.
- [1189] John Robert McDonnell, Robert G. Reynolds, and David B. Fogel, editors. *Proceedings of the 4th Annual Conference on Evolutionary Programming*, A Bradford Book, Complex Adaptive Systems, March 1995, San Diego, CA, USA. The MIT Press, Cambridge, Massachusetts. ISBN 0-262-13317-2.
- [1190] A.V. Sebald and Lawrence J. Fogel, editors. *Proceedings of the 3rd Annual Conference on Evolutionary Programming*, 1994. World Scientific Publishing, River Edge, NJ.
- [1191] David B. Fogel and W. Atmar, editors. *Proceedings of the 2nd Annual Conference on Evolutionary Programming*, February 25–26, 1993, La Jolla, CA, USA. Evolutionary Programming Society, 9363 Towne Centre Dr., San Diego, CA 92121, USA.
- [1192] David B. Fogel and W. Atmar, editors. *Proceedings of the 1st Annual Conference on Evolutionary Programming*, 1992. Evolutionary Programming Society, 9363 Towne Centre Dr., San Diego, CA 92121, USA.
- [1193] John Henry Holland. Adaptive Algorithms for Discovering and Using General Patterns in Growing Knowledge Bases. *International Journal of Policy Analysis and Information Systems*, 4(3):245–268, 1980.
- [1194] Randall Davis and Jonathan King. An Overview of Production Systems. Technical Report STAN-CS-75-524, Stanford Computer Science Department, Stanford University, Stanford Computer Science Department, October 1975. Online available at <http://handle.dtic.mil/100.2/ADA019702> [accessed 2007-07-18].

- [1195] Randall Davis and Jonathan King. An Overview of Production Systems. *Machine Intelligence*, 8, 1977. ISSN 300332.
- [1196] John Henry Holland and Judith S. Reitman. Cognitive systems based on adaptive algorithms. *ACM SIGART Bulletin*, 63:49, June 1977. ISSN 0163-5719. Online available at <http://doi.acm.org/10.1145/1045343.1045373> [accessed 2007-08-18]. See also [1067].
- [1197] Robert Elliott Smith. A Report on The First International Workshop on Learning Classifier Systems (IWLCS-92). Online available at <http://en.scientificcommons.org/70956> and <http://libra.msra.cn/paperDetail.aspx?id=352868> [accessed 2007-08-18]. See [1219], 1992.
- [1198] David E. Moriarty, Alan C. Schultz, and John J. Grefenstette. Evolutionary Algorithms for Reinforcement Learning. *Journal of Artificial Intelligence Research*, 11:241–276, September 1999. Online available at <http://citeseer.ist.psu.edu/moriarty99evolutionary.html> and <http://www.jair.org/media/613/live-613-1809-jair.pdf> [accessed 2007-09-12].
- [1199] Yang Gao, Joshua Zhexue Huang, Hongqiang Rong, and Daqian Gu. Learning classifier system ensemble for data mining. In *GECCO'05: Proceedings of the 2005 workshops on Genetic and evolutionary computation*, pages 63–66, 2005. doi:10.1145/1102256.1102268. In proceedings [555]. Online available at <http://doi.acm.org/10.1145/1102256.1102268> and <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2005wks/papers/0063.pdf> [accessed 2007-09-12].
- [1200] Jaume Bacardit i Peñarroya. *Pittsburgh Genetics-Based Machine Learning in the Data Mining era: Representations, generalization, and run-time*. PhD thesis, Computer Science Department, Enginyeria i Arquitectura La Salle, Ramon Llull University, Barcelona, Catalonia, Spain, October 8, 2004. Advisor: Josep Maria Garrell i Guu. Online available at <http://www.cs.nott.ac.uk/~jqb/publications/thesis.pdf> [accessed 2007-09-12].
- [1201] Hai H. Dam, Hussein A. Abbass, and Chris Lokan. DXCS: an XCS system for distributed data mining. In *Proceedings of Genetic and Evolutionary Computation Conference GECCO 2005*, pages 1883–1890, 2005. In proceedings [554] and also [1938]. Online available at <http://doi.acm.org/10.1145/1068009.1068326> [accessed 2007-09-12].
- [1202] Olgierd Unold and Grzegorz Dabrowski. Use of Learning Classifier System for Inferring Natural Language Grammar. In *Design and Application of Hybrid Intelligent Systems, HIS03, the Third International Conference on Hybrid Intelligent Systems*, pages 272–278, 2003. In proceedings [307] and also [1203].
- [1203] Olgierd Unold and Grzegorz Dabrowski. Use of Learning Classifier System for Inferring Natural Language Grammar. In *Revised Selected Papers of the International Workshops on Learning Classifier Systems*, chapter 2, part I, pages 17–24. Springer Berlin/Heidelberg, 2003. doi:10.1007/978-3-540-71231-2_2. In collection [1212] and also [1202].
- [1204] Walling Cyre. Learning grammars with a modified classifier system. In *CEC'02: Proceedings of the Congress on Evolutionary Computation*, pages 1366–1371, 2002. In proceedings [499].
- [1205] John H. Holmes. Discovering Risk of Disease with a Learning Classifier System. In *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA97)*, 1997. In proceedings [746]. Online available at <http://citeseer.ist.psu.edu/holmes98discovering.html> and <http://www.cs.bris.ac.uk/~kovacs/lcs.archive/Holmes1997a.ps.gz> [accessed 2007-09-12].
- [1206] Jianhua Lin. Adaptive image quantization based on learning classifier systems. In *DCC'95: Proceedings of the Conference on Data Compression*, page 477, March 28–30, 1995. IEEE Computer Society, Washington, DC, USA.

- [1207] A. D. McAulay and Jae C. Oh. Image Learning Classifier System Using Genetic Algorithms. In *Proceedings of IEEE National Aerospace Electronics Conference NAECON'89*, volume 2, pages 705–710, May 22–26, 1989. an IEEE conference.
- [1208] Rick L. Riolo. LETSEQ: An Implementation of the CFS-C Classifier System in a Task-Domain that Involves Learning to Predict Letter. Technical Report, Logic of computers group, Division of computer science and engineering, University of Michigan, 1986, revised 1988. Online available at <http://citeseer.ist.psu.edu/345595.html> [accessed 2007-09-11].
- [1209] *The Tenth International Workshop on Learning Classifier Systems (IW LCS 2007)*, Lecture Notes in Computer Science (LNCS) subseries Lecture Notes in Artificial Intelligence series (LNAI), July 8, 2007, University College London, in London, England. Springer. Held in association with GECCO-2007 (see [551, 552]).
- [1210] *Ninth International Workshop on Learning Classifier Systems (IW LCS 2006)*, July 8–9, 2006, Seattle, WA, USA. Held during GECCO-2006 (see [553]).
- [1211] *Eighth International Workshop on Learning Classifier Systems (IW LCS-2005)*, June 25–29, 2005, Washington DC, USA. Held during GECCO-2005 (see [554, 555]). See also [1212].
- [1212] Tim Kovacs, Xavier Llorà, Keiki Takadama, Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors. *Revised Selected Papers of the International Workshops on Learning Classifier Systems, IW LCS 2003-2005*, volume 4399/2007 of *Lecture Notes in Computer Science (LNCS)*, subseries *Lecture Notes in Artificial Intelligence (LNAI)*, April 19, 2007. Springer Berlin/Heidelberg. ISBN 978-3-540-71230-5. doi:10.1007/978-3-540-71231-2.
- [1213] *Seventh International Workshop on Learning Classifier Systems (IW LCS-2004)*, June 26, 2004, Seattle, Washington, USA. Held during GECCO-2004 (see [558, 559]). See also [1212].
- [1214] *Sixth International Workshop on Learning Classifier Systems (IW LCS-2003)*, July 12–16, 2003, The Holiday Inn Chicago, Chicago, IL 60654, USA. Held during GECCO-2003 (see [561, 562]). See also [1212].
- [1215] Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors. *Revised Papers of the 5th International Workshop on Learning Classifier Systems, IW LCS 2002*, volume 2661/2003 of *Lecture Notes in Computer Science (LNCS)*, September 7–8, 2002, Granada, Spain. Springer. ISBN 3-540-20544-6. Published in 2003.
- [1216] *Advances in Learning Classifier Systems, Revised Papers of the Fourth International Workshop on Learning Classifier Systems (IW LCS-2001)*, volume 2321/2002 of *Lecture Notes in Computer Science (LNCS)*, July 7–8, 2001, San Francisco, CA, USA. Springer. ISBN 3-540-43793-2. Held during GECCO-2001 (see [568]), published in 2002.
- [1217] Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors. *Advances in Learning Classifier Systems, Revised Papers of the Third International Workshop on Learning Classifier Systems (IW LCS-2000)*, volume 1996/2001 of *Lecture Notes in Computer Science (LNCS)*, September 15–16, 2000, Paris, France. Springer. ISBN 3-540-42437-7. A joint workshop of the sixth International Conference on Simulation of Adaptive Behaviour (SAB2000, see also [1939]) and the sixth International Conference on Parallel Problem Solving from Nature (PPSN VI, see also [587]), published in 2001.
- [1218] *The Second International Workshop on Learning Classifier Systems (IW LCS-99)*, July 13, 1999, Orlando, Florida, USA. Co-located with GECCO-99 (see [572]). Proceedings are published in [1224].
- [1219] *Collected Abstracts for the First International Workshop on Learning Classifier Systems (IW LCS-92)*, October 6–9, 1992, NASA Johnson Space Center in Houston, Texas, USA.

- [1220] Larry Bull, editor. *Applications of Learning Classifier Systems*. Studies in Fuzziness and Soft Computing. Springer, Berlin, Germany, April 2004. ISBN 3540211098, 978-3540211099. Partly online available at <http://books.google.de/books?id=aBIjqGag5-kC> [accessed 2008-04-01].
- [1221] Larry Bull and Tim Kovacs, editors. *Foundations of Learning Classifier Systems*. Studies in Fuzziness and Soft Computing. Springer, September 1, 2005. ISBN 3540250735, 978-3540250739.
- [1222] Martin V. Butz. *Anticipatory Learning Classifier Systems*. Genetic Algorithms and Evolutionary Computation. Springer, January 25 2002. ISBN 0792376307, 978-0792376309.
- [1223] Martin V. Butz. *Rule-Based Evolutionary Online Learning Systems: A Principled Approach to LCS Analysis and Design*. Studies in Fuzziness and Soft Computing. Springer, December 22 2005. ISBN 3540253793, 978-3540253792.
- [1224] Pier Luca Lanzi, Wolfgang Stolzmann, and Steward W. Wilson, editors. *Learning Classifier Systems, From Foundations to Applications*, volume 1813/2000 of *Lecture Notes in Artificial Intelligence, subseries of Lecture Notes in Computer Science (LNCS)*. Springer-Verlag Berlin/Heidelberg, London, UK, 2000. ISBN 3-540-67729-1.
- [1225] Andreas Geyer-Schulz from the Department of Applied Computer Science, Vienna University of Economics and Business Administration, A-1090 Vienna, Augasse 2-6, Austria. Holland classifier systems. In *APL'95: Proceedings of the international conference on Applied programming languages*, pages 43–55, 1995, San Antonio, Texas, United States. ACM Press, New York, NY, USA. ISBN 0-89791-722-7. doi:10.1145/206913.206955. Online available at <http://doi.acm.org/10.1145/206913.206955> [accessed 2007-09-12].
- [1226] Bart de Boer. Classifier Systems: a useful approach to machine learning? Master's thesis, Leiden University, Rijksuniversiteit Leiden, Netherlands, August 1994. Internal Report Number IR94-02. Supervisors Ida Sprinkhuizen-Kuyper and Egbert Boers. Online available at citeseer.ist.psu.edu/deboer94classifier.html [accessed 2007-08-08].
- [1227] Marvin Lee Minsky. *Berechnung: Endliche und Unendliche Maschinen*. Verlag Berliner Union GmbH, Stuttgart, 1971. ISBN 978-3408530294.
- [1228] Marvin Lee Minsky. *Computation: Finite and Infinite Machines*. Prentice Hall Series on Automatic Computation. Prentice Hall, June 1967. ISBN 978-0131655638.
- [1229] Tommaso Toffoli. Reversible Computing. In *Proceedings of the 7th Colloquium on Automata, Languages and Programming*, Lecture Notes In Computer Science (LNCS), pages 632–644, 1980. Springer-Verlag, London, UK. ISBN 3-540-10003-2.
- [1230] John Henry Holland, Keith J. Holyoak, and Richard E. Nisbett and Paul R. Thagard. *Induction: Processes of Inference, Learning, and Discovery*. Computational Models of Cognition & Perception. The MIT Press, first, 1989 reprint edition, March 1986. ISBN 0262580969, 978-0262580960. Partly online available at <http://books.google.de/books?id=Z6EFBaLApE8C> [accessed 2008-04-03].
- [1231] Rick L. Riolo. Bucket brigade performance: II. Default hierarchies. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 196–201, 1987. In proceedings [750].
- [1232] Rick L. Riolo. The Emergence of Default Hierarchies in Learning Classifier Systems. In *Proceedings of the 3rd International Conference on Genetic Algorithms (ICGA89)*, pages 322–327, 1989. In proceedings [652].
- [1233] Robert Elliott Smith. *Default hierarchy formation and memory exploitation in learning classifier systems*. PhD thesis, University of Alabama, Tuscaloosa, AL, USA, 1991. UMI Order No. GAX91-30265.
- [1234] John Henry Holland. Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning: An Artificial Intel-*

- ligence Approach: Volume II*, pages 593–623. Kaufmann, Los Altos, CA, 1986. See also [1940].
- [1235] Marvin Lee Minsky. Steps toward artificial intelligence. In *Proceedings of the IRE*, volume 49(1), pages 8–30, 1961. Reprint: Feigenbaum & Feldman, Computers and Thought, 1963, Iuger.
 - [1236] Marvin Lee Minsky. Steps toward artificial intelligence. In Edward A. Feigenbaum and Julian Feldman, editors, *Computers & thought*, pages 406–450. MIT Press, Cambridge, MA, USA, 1995. ISBN 0-262-56092-5. Online available at <http://web.media.mit.edu/~minsky/papers/steps.html> [accessed 2007-08-05].
 - [1237] John Henry Holland. Properties of the Bucket Brigade Algorithm. In *Proceedings of the International Conference on Genetic Algorithms and Their Applications*, pages 1–7, 1985. In proceedings [751].
 - [1238] Thomas H. Westerdale. The Bucket Brigade Is Not Genetic. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pages 45–59, 1985. In proceedings [751].
 - [1239] Thomas H. Westerdale. A reward scheme for production systems with overlapping conflict sets. *IEEE Transactions on Systems, Man and Cybernetics*, 16(3):369–383, 1986. ISSN 0018-9472.
 - [1240] Thomas H. Westerdale. Altruism in the bucket brigade. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 22–26, 1987. In proceedings [750].
 - [1241] Hendrik James Antonisse. Unsupervised credit assignment in knowledge-based sensor fusionsystems. *IEEE Transactions on Systems, Man and Cybernetics*, 20(5):1153–1171, September/October 1990. doi:10.1109/21.59978.
 - [1242] D. Huang. A framework for the credit-apportionment process in rule-based systems. *IEEE Transactions on Systems, Man and Cybernetics*, 19(3):489–498, May/June 1989. doi:10.1109/21.31056.
 - [1243] Rick L. Riolo. Bucket brigade performance: I. Long sequences of classifiers. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 184–195, 1987. In proceedings [750].
 - [1244] Marco Dorigo. Message-Based Bucket Brigade: An Algorithm for the Apportionment of Credit Problem. In Yves Kodratoff, editor, *EWSL-91: Proceedings of the European working session on learning on Machine learning*, pages 235–244, 1991, Porto, Portugal. Springer-Verlag New York, Inc., New York, NY, USA. ISBN 0-387-53816-X. Number 482. Online available at <http://citeseer.ist.psu.edu/dorigo91messagebased.html> [accessed 2008-04-04].
 - [1245] Piet Spiessens. PCS: A classifier system that builds a predictive internal world model. In Luigia Carlucci Aiello, editor, *ECAI 90 9th European Conference on Artificial Intelligence*, pages 622–627, August 6–10 1990, Stockholm, Sweden. Pitman Publishing, London, UK.
 - [1246] Stewart Wilson. Bid competition and specificity reconsidered. *Complex Systems*, 2(6):705–723, 1988. ISSN 0891-2513.
 - [1247] Nabil M. Hewahi and K. K. Bharadwaj. Bucket brigade algorithm for hierarchical censored production rule-based system. *International Journal of Intelligent Systems*, 11(4):197–225, 1996.
 - [1248] Nabil M. Hewahi from the Department of Computer Science, The Islamic University of Gaza, Palestine, P.O. Box 108, Gaza, Palestine. Credit Apportionment Scheme for Rule Based Systems. *Journal of the Islamic University of Gaza*, 9(1):25–41, 2001.
 - [1249] Nabil M. Hewahi and H. Ahmad. Credit Apportionment Scheme for Rule-Based Systems: Implementation and Comparative Study. In *Developments in Applied Artificial Intelligence, Proceedings of the 15th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems IEA/AIE*, volume 2358/2002 of *Lecture Notes in Computer Science (LNCS)*, pages 43–56, June 17–

- 20, 2002, Cairns, Australia. Springer, Berlin/Heidelberg. ISBN 978-3-540-43781-9. doi:10.1007/3-540-48035-8_43.
- [1250] Stewart W. Wilson. ZCS: A Zeroth Level Classifier System. *Evolutionary Computation*, 2(1):1–18, 1994. Online available at <http://citeseer.ist.psu.edu/wilson94zcs.html> and <http://www.eskimo.com/~wilson/ps/zcs.pdf> [accessed 2007-09-12].
- [1251] Marco Dorigo and Hugues Bersini. A Comparison of Q-Learning and Classifier Systems. In *Proceedings of From Animals to Animats, Third International Conference on Simulation of Adaptive Behavior*, pages 248–255, 1994, Brighton, United Kingdom. ISBN 0-262-53122-4. Online available at <http://citeseer.ist.psu.edu/dorigo94comparison.html> [accessed 2008-04-03].
- [1252] Tim Kovacs. Two Views of Classifier Systems. In *Fourth International Workshop on Learning Classifier Systems - IW LCS-2001*, pages 367–371, 7 2001. San Francisco, California, USA. In proceedings [1216]. Online available at <http://citeseer.ist.psu.edu/kovacs02two.html> and <http://www.cs.bris.ac.uk/Publications/Papers/1000647.pdf> [accessed 2007-09-11].
- [1253] John Henry Holland, Lashon Bernard Booker, Marco Colombetti, Marco Dorigo, David E. Goldberg, Stephanie Forrest, Rick L. Riolo, Robert Elliott Smith, Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson. What Is a Learning Classifier System? In *Learning Classifier Systems, From Foundations to Applications*, pages 3–32. Springer-Verlag Berlin/Heidelberg, 2000. In collection [1224].
- [1254] Pier Luca Lanzi and Rick L. Riolo. A Roadmap to the Last Decade of Learning Classifier System Research (From 1989 to 1999). In *Learning Classifier Systems: From Foundations to Applications*, page 33. Springer-Verlag Berlin/Heidelberg, 2000. In collection [1224].
- [1255] Stewart W. Wilson and David E. Goldberg. A Critical Review of Classifier Systems. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 244–255, 1989. In proceedings [652]. Online available at <http://www.eskimo.com/~wilson/ps/wg.ps.gz> [accessed 2007-09-12].
- [1256] H. Brown Cribbs, III and Robert Elliott Smith. What Can I Do With a Learning Classifier System? In *Industrial Applications of Genetic Algorithms*, pages 299–319. CRC Press, 1999. In collection [764].
- [1257] Jason Brownlee. Learning Classifier Systems. Technical Report 070514A, Complex Intelligent Systems Laboratory, Centre for Information Technology Research, Faculty of Information and Communication Technologies, Swinburne University of Technology Melbourne, Australia, May 2007. Online available at <http://www.ict.swin.edu.au/personal/jbrownlee/2007/TR24-2007.pdf> [accessed 2008-04-03].
- [1258] Kenneth Alan De Jong. Learning with genetic algorithms: An overview. *Machine Learning*, 3(2-3):121–138, October 1988. ISSN 0885-6125 (Print) 1573-0565 (Online). doi:10.1007/BF00113894. Online available at <http://www.springerlink.com/content/mg877066r521w781/fulltext.pdf> [accessed 2008-04-04].
- [1259] Kenneth Alan De Jong. On using genetic algorithms to search program spaces. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 210–216, 1987. In proceedings [750].
- [1260] John J. Grefenstette from the Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, Washington, DC 20375-5000, U.S.A. Credit Assignment in Rule Discovery Systems Based on Genetic Algorithms. *Zeitschrift Machine Learning*, 3:225–245, October 1988. ISSN 0885-6125 (Print) 1573-0565 (Online). doi:10.1023/A:1022614421909. Received: December 21, 1987, Revised: July 15, 1988. Online available at <http://www.springerlink.com/content/n83r83v087685215/fulltext.pdf> [accessed 2008-05-29].
- [1261] Kenneth Alan De Jong and William M. Spears. Learning Concept Classification Rules using Genetic Algorithms. In *Proceedings of the Twelfth International Conference*

- on *Artificial Intelligence IJCAI-91*, volume 2, 1991. In proceedings [1941]. Online available at <http://citeseer.ist.psu.edu/dejong91learning.html> and <http://www.cs.uwo.edu/~wspears/papers/ijcai91.pdf> [accessed 2007-09-12].
- [1262] Kenneth Alan De Jong, William M. Spears, and Diana F. Gordon. Using Genetic Algorithms for Concept Learning. *Machine Learning*, 13:161–188, 1993. Online available at <http://citeseer.ist.psu.edu/dejong93using.html> and <http://www.cs.uwo.edu/~dspears/papers/mlj93.pdf> [accessed 2007-09-12].
- [1263] Jaume Bacardit i Peñarroya. Analysis of the initialization stage of a Pittsburgh approach learning classifier system. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1843–1850, 2005. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068321> and <http://www.cs.nott.ac.uk/~jqb/publications/gecco2005.pdf> [accessed 2007-09-12].
- [1264] Jaume Bacardit i Peñarroya and Natalio Krasnogor. Smart crossover operator with multiple parents for a Pittsburgh learning classifier system. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1441–1448, 2006. In proceedings [553]. Online available at <http://doi.acm.org/10.1145/1143997.1144235> and <http://www.cs.nott.ac.uk/~jqb/publications/gecco2006-sx.pdf> [accessed 2007-09-12].
- [1265] David E. Goldberg. *Computer-aided gas pipeline operation using genetic algorithms and rule learning*. PhD thesis, University of Michigan. Ann Arbor, MI, 1983.
- [1266] Xavier Llorà and Kumara Sastry. Fast rule matching for learning classifier systems via vector instructions. In *GECCO'06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1513–1520, 2006. doi:10.1145/1143997.1144244. In proceedings [553] and also [1942]. Online available at <http://doi.acm.org/10.1145/1143997.1144244> [accessed 2007-09-12].
- [1267] Stewart W. Wilson. Classifier Fitness Based on Accuracy. *Evolutionary Computation*, 3(2):149–175, 1995. Online available at <http://citeseer.ist.psu.edu/wilson95classifier.html> [accessed 2007-09-12].
- [1268] Dave Cliff and Susi Ross. Adding temporary memory to ZCS. *Adaptive Behavior*, 3(2):101–150, Fall 1994. ISSN 1059-7123. Online available at <ftp://ftp.informatics.sussex.ac.uk/pub/reports/csrp/csrp347.ps.Z> [accessed 2007-09-12].
- [1269] Larry Bull and Jacob Hurst. ZCS redux. *Evolutionary Computation*, 10(2):185–205, 2002. ISSN 1063-6560. Online available at <http://www.cems.uwe.ac.uk/lcsg/papers/zcsredux.ps> [accessed 2007-09-12].
- [1270] Larry Bull. On using ZCS in a Simulated Continuous Double-Auction Market. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)*, pages 83–90, 1999. In proceedings [572]. Online available at <http://www.cs.bham.ac.uk/~wbl/biblio/gecco1999/GA-806.pdf> [accessed 2007-09-12].
- [1271] Tim Kovacs. XCS Classifier System Reliably Evolves Accurate, Complete, and Minimal Representations for Boolean Functions. In *Soft Computing in Engineering Design and Manufacturing*, pages 59–68. Springer-Verlag, 1998. In collection [271]. Online available at <http://www.cs.bris.ac.uk/Publications/Papers/1000239.pdf> [accessed 2007-08-18].
- [1272] Stewart W. Wilson. Generalization in the XCS Classifier System. In *Proceedings of the Third Annual Conference on Genetic Programming 1998*, pages 665–674, 1998. In proceedings [913]. Online available at <http://citeseer.ist.psu.edu/wilson98generalization.html> [accessed 2007-09-12].
- [1273] Stewart W. Wilson. State of XCS Classifier System Research. In *Learning Classifier Systems, From Foundations to Applications*, pages 63–82. Springer-Verlag, 2000. In collection [1224]. Online available at <http://citeseer.ist.psu.edu/72750.html> and <http://www.eskimo.com/~wilson/ps/state.ps.gz> [accessed 2007-08-23].

- [1274] Pier Luca Lanzi, Daniele Loiacono, Stewart W. Wilson, and David E. Goldberg. Generalization in the XCSF Classifier System: Analysis, Improvement, and Extension. *Evolutionary Computation Journal*, 2006. See also [1943].
- [1275] Tash. Body Rock. In Mos Def, Q-Tip, and Tash, editors, *Body Rock (Single)*. Rawkus Entertainment, March 3, 1998. ASIN: B000001XRQ.
- [1276] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, second edition, December 2002. ISBN 0137903952.
- [1277] Bowei Xi, Zhen Liu, Mukund Raghavachari, Cathy H. Xia, and Li Zhang. A smart hill-climbing algorithm for application server configuration. In *WWW'04: Proceedings of the 13th international conference on World Wide Web*, pages 287–296, 2004, New York, NY, USA. ACM Press, New York, NY, USA. ISBN 1-58113-844-X. doi:10.1145/988672.988711. Online available at <http://citeseer.ist.psu.edu/xi04smart.html> [accessed 2007-09-11].
- [1278] Brian P. Gerkey, Sebastian Thrun, and Geoff Gordon. Parallel stochastic hill-climbing with small teams. In Lynne E. Parker, Frank E. Schneider, and Alan C. Schultz, editors, *Proceedings from the 2005 International Workshop on Multi-Robot Systems*, volume III of *Multi-Robot Systems: From Swarms to Intelligent Automata*, pages 65–77. Springer, 2005. ISBN 978-1-4020-3388-9. Presented at 2005 International Workshop on Multi-Robot Systems. L.E.Parker et al. Online available at <http://www.cs.cmu.edu/~ggordon/gerkey-thrun-gordon.parish.pdf> [accessed 2007-08-18].
- [1279] Michael E. Farmer, Shweta Bapna, and Anil K. Jain. Large Scale Feature Selection Using Modified Random Mutation Hill Climbing. In *ICPR 2004. Proceedings of the 17th International Conference on Pattern Recognition*, volume 2, pages 287–290, August 23–26, 2004. IEEE Computer Society, Los Alamitos, CA, USA. doi:10.1109/ICPR.2004.1334169.
- [1280] Alexandre Temporel and Tim Kovacs from the Department of Computer Science, University of Bristol. A heuristic hill climbing algorithm for mastermind. In *Proceedings of the 2003 UK Workshop on Computational Intelligence (UKCI-03)*, pages 189–196. University of Bristol, September 1–3, 2003, Department of Engineering Mathematics and Department of Computer Science The University of Bristol, UK. ISBN 0862925371. Online available at <http://citeseer.ist.psu.edu/701434.html> and <http://www.cs.bris.ac.uk/Publications/Papers/2000067.pdf> [accessed 2007-09-11].
- [1281] Robert C. Holte from the School of Information Technology and Engineering, University of Ottawa, Ottawa, Ontario, Canada, K1N 6N5. Combinatorial Auctions, Knapsack Problems, and Hill-Climbing Search. In *Proceedings of the 14th Biennial Conference of the Canadian Society on Computational Studies of Intelligence: Advances in Artificial Intelligence*, volume 2056/2001 of *Lecture Notes in Computer Science (LNCS)*, page 57. Springer Berlin / Heidelberg, June 7–9, 2001, Ottawa, Canada. ISBN 3-540-42144-0. Online available at <http://citeseer.ist.psu.edu/holte01combinatorial.html> [accessed 2007-09-11].
- [1282] Ted Carson and Russell Impagliazzo. Hill-climbing finds random planted bisections. In *Symposium on Discrete Algorithms, Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, volume 12, pages 903–909. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001, Washington, D.C., United States. ISBN 0-89871-490-7. Online available at <http://portal.acm.org/citation.cfm?id=365411.365805> and <http://citeseer.ist.psu.edu/carson01hillclimbing.html> [accessed 2007-09-11].
- [1283] Deniz Yuret and Michael de la Maza. Dynamic hill climbing: Overcoming the limitations of optimization techniques. In *Proceedings of the Second Turkish Symposium on Artificial Intelligence and Neural Networks*, pages 208–212, June 24–25, 1993, Boğaziçi University, Istanbul, Turkey. Online available at <http://citeseer.ist.psu.edu/yuret93dynamic.html> and <http://www.denizyuret.com/pub/tainn93.html> [accessed 2007-09-11].

- [1284] James C. Spall. *Introduction to Stochastic Search and Optimization*. Estimation, Simulation, and Control – Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, first edition, June 2003. ISBN 978-0-471-33052-3.
- [1285] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. Estimation, Simulation, and Control – Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley Interscience, second edition, November 2000. ISBN 0-471-05669-3.
- [1286] Leonard A. Rastrigin. The Convergence of the Random Search Method in the Extremal Control of Many-Parameter System. *Automation and Remote Control*, 24: 1337–1342, 1963.
- [1287] Michael A. Schumer. *Optimization by adaptive random search*. PhD thesis, Princeton University, NJ, November 1967. Supervisor Kenneth Steiglitz.
- [1288] Michael A. Schumer and Kenneth Steiglitz. Adaptive step size random search. *IEEE Transactions on Automatic Control*, AC-13(3):270–276, June 1968.
- [1289] R. C. White, jr from the Department of Electrical Engineering Technological University Eindhoven, Netherlands. A survey of random methods for parameter optimization. *Simulation*, 17(5):197–205, 1971. doi:10.1177/003754977101700504. Online available at <http://sim.sagepub.com/cgi/reprint/17/5/197?ck=nck> [accessed 2008-03-26].
- [1290] J. Matyas. Random Optimization. *Automation and Remote Control*, 26(2):244–251, 1965.
- [1291] J. Matyas. Das Zufällige Optimierungsverfahren und Seine Konvergenz. In *Proceedings of the 5th International Analogue Computation Meeting*, pages 540–544. Presses Academiques Europeens, Brussels, 1968, Lausanne.
- [1292] N. Baba from the Information Science and Systems Engineering, Faculty of Engineering, Tokushima University, Tokushima City, Japan. Convergence of a random optimization method for constrained optimization problems. *Journal of Optimization Theory and Applications*, 33(4):451–461, April 1981. ISSN 0022-3239 (Print) 1573-2878 (Online). Communicated by N. Avriel, Online available at <http://www.springerlink.com/content/q1g0k627688684k3/> [accessed 2007-08-26].
- [1293] Nimit Worakul, Wibul Wongpoowarak, and Prapaporn Boonme. Optimization in Development of Acetaminophen Syrup Formulation. *Drug Development and Industrial Pharmacy*, 28(3):345–351, 2002. ISSN 1520-5762 (electronic) 0363-9045 (paper). see erratum [1294].
- [1294] Nimit Worakul, Wibul Wongpoowarak, and Prapaporn Boonme. Optimization in Development of Acetaminophen Syrup Formulation – Erratum. *Drug Development and Industrial Pharmacy*, 28(8):1043–10045, 2002. ISSN 1520-5762 (electronic) 0363-9045 (paper). see paper [1293].
- [1295] Junyi Li and R. Russell Rhinehart from the Department of Chemical Engineering, Texas Tech University, Lubbock, TX 79409, U.S.A. Heuristic random optimization. *Computers and Chemical Engineering*, 22(3):427–444, February 1998. doi:10.1016/S0098-1354(97)00005-7. Received 10 August 1995; revised 28 March 1996.
- [1296] Sandeep Chandran and R. Russell Rhinehart. Heuristic random optimizer-version II. In *Proceedings of the American Control Conference*, volume 4, pages 2589–2594. IEEE, Piscataway NJ, US, 2002. doi:10.1109/ACC.2002.1025175.
- [1297] Sanjeev Dhir, K. John Morrow Jr, R. Russell Rhinehart, and Theodore Wiesner. Dynamic optimization of hybridoma growth in a fed-batch bioreactor. *Biotechnology and Bioengineering*, 67(2):197–205, January 2000. ISSN 1097-0290. doi:10.1002/(SICI)1097-0290(20000120)67:2<197::AID-BIT9>3.0.CO;2-W. Online available at <http://www3.interscience.wiley.com/cgi-bin/abstract/71003651/> [accessed 2007-08-26].

- [1298] Pedro A. González Lanza and Jesús M. Zamarreño Cosme. A short-term temperature forecaster based on a state space neural network. *Engineering Applications of Artificial Intelligence*, 15(5):459–464, September 2002. doi:10.1016/S0952-1976(02)00089-1.
- [1299] Francisco J. Solis and Roger J-B. Wets. Minimization by random search techniques. *Mathematics of operations research*, 6(1):19–30, February 1981. ISSN 0364765X.
- [1300] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, June 1953. doi:10.1063/1.1699114. Online available at <http://link.aip.org/link/?JCPSA6/21/1087/1> [accessed 2008-03-26].
- [1301] Scott Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598):671–680, May 13, 1983. Online available at <http://fezzik.ucd.ie/msc/cscs/ga/kirkpatrick83optimization.pdf> and <http://citeseer.ist.psu.edu/kirkpatrick83optimization.html> [accessed 2008-03-26].
- [1302] V. Černý from the Institute of Physics and Biophysics, Comenius University, Mlynska Dolina, Bratislava, Czechoslovakia. Thermodynamical Approach to the Traveling Salesman Problem: An Efficient Simulation Algorithm. *Journal of Optimization Theory and Applications*, 45(1):41–51, January 1985. ISSN 0022-3239 (Print) 1573-2878 (Online). doi:10.1007/BF00940812. Communicated by S. E. Dreyfus. Online available at <http://www.springerlink.com/content/gt0743622913gg33/fulltext.pdf> [accessed 2008-03-26]. Also: Technical Report, Comenius University, Bratislava, Czechoslovakia, 1982.
- [1303] Andreas Nolte and Rainer Schrader from the Mathematisches Institut, Universität zu Köln. A Note on the Finite Time Behaviour of Simulated Annealing. *Mathematics of OR*, 25(347):476–484, 2000. Online available at <http://citeseer.ist.psu.edu/nolte99note.html> and <http://www.zaik.de/~paper/unzip.html?file=zaik1999-347.ps> [accessed 2008-03-26] (revised version from March 1999). See also [1308].
- [1304] P. J. M. Laarhoven and E. H. L. Aarts, editors. *Simulated Annealing: Theory and Applications*. Mathematics and Its Applications. Springer, Kluwer Academic Publishers, Norwell, MA, USA, June 30, 1987. ISBN 9027725136, 978-9027725134, 31289400, 40292200. Reviewed in [1944]. Partly online available at http://books.google.de/books?id=-IgUab6Dp_IC [accessed 2008-03-26].
- [1305] B. Hajek. Cooling Schedules for Optimal Annealing. *Mathematics of Operations Research*, (13), 1988.
- [1306] S. Anily and Awi Federgruen. Ergodicity in parametric non-stationary Markov chains: An Application to Simulated Annealing Methods. *Operations Research*, (35):867–874, 1987.
- [1307] Basilis Gidas. Nonstationary Markov chains and convergence of the annealing algorithm. *Journal of Statistical Physics*, 39(1-2):73–131, April 1985. ISSN 0022-4715 (Print) 1572-9613 (Online). doi:10.1007/BF01007975. Online available at <http://www.springerlink.com/content/x3821j328gv33085/fulltext.pdf> [accessed 2008-03-26].
- [1308] Andreas Nolte and Rainer Schrader. A Note on the Finite Time Behaviour of Simulated Annealing. In U. Zimmermann, U. Derigs, W. Gaul, R. Möhring, and K.-P. Schuster, editors, *Operations Research Proceedings 1996, Selected Papers of the Symposium on Operations Research (SOR 96)*, pages 175–180. Springer, 1996. Published 1997. See also [1303].
- [1309] Debasis Mitra, Fabio Romeo, and Alberto Sangiovanni-Vincentelli from the AT & T Bell Laboratories, Murray Hill, NJ. Convergence and finite-time behavior of simulated annealing. In *Proceedings of the 24th IEEE Conference on Decision*

- and *Control*, volume 24, pages 761–767. IEEE Computer Society, December 1985. doi:10.1109/CDC.1985.268600.
- [1310] Lester Ingber. Simulated annealing: Practice versus theory. *Mathematical and Computer Modelling*, 18(11):29–57, 1993. Online available at <http://citeseer.ist.psu.edu/589471.html> [accessed 2007-09-15]. See also [1945].
- [1311] Lester Ingber. Adaptive simulated annealing (ASA): Lessons learned. *Control and Cybernetics*, 25(1):33–54, 1996. Online available at http://www.ingber.com/asa96_lessons.pdf and <http://citeseer.ist.psu.edu/592447.html> [accessed 2008-03-27].
- [1312] Shinji Sato. Simulated quenching: a new placement method for module generation. In *ICCAD'97: Proceedings of the 1997 IEEE/ACM international conference on Computer-aided design*, pages 538–541, 1997, San Jose, California, United States. IEEE Computer Society, Washington, DC, USA. ISBN 0-8186-8200-0. doi:10.1109/ICCAD.1997.643591. Online available at http://www.sigda.org/Archives/ProceedingArchives/Iccad/Iccad97/papers/1997/iccad97/pdffiles/09c_2.pdf and http://portal.acm.org/ft_gateway.cfm?id=266547&type=pdf&coll=GUIDE&dl=GUIDE&CFID=61186331&CFTOKEN=74233409 [accessed 2008-03-27].
- [1313] Jean-Numa Gillet and Yunlong Sheng. Simulated quenching with temperature rescaling for designing diffractive optical elements. In Alexander J. Glass, Joseph W. Goodman, Milton Chang, Arthur H. Guenther, and Toshimitsu Asakura, editors, *Proceedings of the 18th Congress of the International Commission for Optics*, volume 3749 of *Proceedings of SPIE*, pages 683–684, July 1999.
- [1314] Bradley J. Buckham and Casey Lambert from the Mechanical Engineering Department, University of Victoria. Simulated Annealing Applications, November 1999. Seminar presentation: MECH620 Quantitative Analysis, Reasoning and Optimization Methods in CAD/CAM and Concurrent Engineering, Department of Mechanical Engineering, University of Victoria, Course Homepage: <http://www.me.uvic.ca/~zdong/courses/mech620/> [accessed 2007-08-25], instructor: Dr. Zuomin Dong, Online available at http://www.me.uvic.ca/~zdong/courses/mech620/SA_App.PDF [accessed 2007-08-25].
- [1315] William L. Goffe, Gary D. Ferrier, and John Rogers. Global Optimization of Statistical Functions With Simulated Annealing. *Journal of Econometrics*, 60(1-2):65–99, January-February 1994. Online available at [http://dx.doi.org/10.1016/0304-4076\(94\)90038-8](http://dx.doi.org/10.1016/0304-4076(94)90038-8) and <http://cook.rfe.org/anneal-preprint.pdf> [accessed 2007-09-15]. See also <http://econpapers.repec.org/software/codfortra/simanneal.htm> [accessed 2007-09-15].
- [1316] Daniel A. Liotard. Algorithmic Tools in the Study of Semiempirical Potential Surfaces. *International Journal of Quantum Chemistry*, 44(5):723–741, November 5, 1992. doi:10.1002/qua.560440505. Online available at <http://www3.interscience.wiley.com/cgi-bin/fulltext/109573913/PDFSTART> [accessed 2007-09-15].
- [1317] Axel T. Brünger, Paul D. Adams, and Luke M. Rice from the The Howard Hughes Medical Institute and Department of Molecular Biophysics and Biochemistry, Yale University, New Haven, CT 06520, USA. New applications of simulated annealing in X-ray crystallography and solution NMR. *Structure*, 15:325–336, 1997. Online available at <http://atb.slac.stanford.edu/public/papers.php?sendfile=44> [accessed 2007-08-25].
- [1318] Erik Sundermann and Ignace L. Lemahieu. PET image reconstruction using simulated annealing. In Murray H. Loew, editor, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference, Medical Imaging 1995: Image Processing*, volume 2434, pages 378–386, May 1995. doi:10.1117/12.208709.
- [1319] Koon-Pong Wong, S. R. Meikle, Dagan Feng, and M. J. Fulham. Estimation of input function and kinetic parameters using simulated annealing: application in a

- flow model. *IEEE Transactions on Nuclear Science*, 49(3):707–713, June 2002. ISSN 0018-9499. doi:10.1109/TNS.2002.1039552.
- [1320] Maqsood Yaqub, Ronald Boellaard, Marc A Kropholler, and Adriaan A Lammertsma from the Department of Nuclear Medicine & PET Research, VU University Medical Centre, Amsterdam, The Netherlands. Optimization algorithms and weighting factors for analysis of dynamic PET studies. *Physics in Medicine and Biology*, 51:4217–4232, August 2006. doi:10.1088/0031-9155/51/17/007. Online available at stacks.iop.org/PMB/51/4217 [accessed 2007-08-25].
 - [1321] Lester Ingber from the DRW Investments LLC, Chicago Mercantile Exchange Center. A simple options training model. *Mathematical Computer Modelling*, 30(5–6): 167–182, 1999. Online available at http://www.ingber.com/markets99_spread.pdf and <http://citeseer.ist.psu.edu/ingber99simple.html> [accessed 2007-08-25].
 - [1322] Lester Ingber. Trading Markets With Canonical Momenta and Adaptive Simulated Annealing. Technical Report 1996:TMCMSA, Lester Ingber Research, McLean, VA, 1996. Online available at http://www.ingber.com/markets96_brief.pdf and <http://www.geocities.com/francorbusetti/ingbergermarkets.pdf> [accessed 2007-08-25].
 - [1323] R. A. Rutenbar. Simulated annealing algorithms: an overview. *IEEE Circuits and Devices Magazine*, 5(1):19–26, January 1989. doi:10.1109/101.17235.
 - [1324] H. Martinez-Alfaro and D. R. Flugrad. Collision-free path planning for mobile robots and/or AGVs using simulated annealing. In *Systems, Man, and Cybernetics 1994. Proceedings of the IEEE International Conference on Humans, Information and Technology*, volume 1, pages 270–275. IEEE, October 1994. doi:10.1109/ICSMC.1994.399849.
 - [1325] A. Malhotra, J. H. Oliver, and W. Tu. Synthesis of Spatially and Intrinsically Constrained Curves Using Simulated Annealing. *ASME Transactions, Journal of Mechanical Design*, 118:53–61, March 1996. DE-vol. 32-1 Alliances in Design Automation, vol. 1 ASME 1991 pp. 145–155.
 - [1326] I. Ullah and S. Kota. Globally-Optimal Synthesis of Mechanisms for Path Generation Using Simulated Annealing and Powell’s Method. In *Proceedings of the 1996 ASME Design Engineering Technical Conferences and Computers in Engineering*, page 8, August 1996, Irvine/CA, USA.
 - [1327] H. Martinez-Alfaro, H. Valdez, and J. Ortega. Linkage Synthesis of a four bar Mechanism for N Precision Points Using Simulated Annealing. In *Proceedings of the 1998 ASME Design Engineering Technical Conference*, page 7, September 13–16, 1998, Atlanta, USA.
 - [1328] Horacio Martínez-Alfaro and Manuel Valenzuela-Rendón from the ITESM, Center for Intelligent Systems Monterrey, N.L. 64849 México. Using Simulated Annealing for Paper Cutting Optimization. In *MICAI 2004: Advances in Artificial Intelligence*, pages 11–20, 2004. In proceedings [396].
 - [1329] Mrinal K. Sen and Paul L. Stoffa. Nonlinear One-Dimensional Seismic Waveform Inversion Using Simulated Annealing. *Geophysics*, 56(10):1624–1638, October 1991.
 - [1330] Adedeji Bodunde Badiru. *Handbook of Industrial and Systems Engineering*. CRC Press Inc, January 2006. ISBN 0849327199, 978-0849327193. Partly online available at <http://books.google.de/books?id=4FkmK2A6wuUC> [accessed 2008-03-27].
 - [1331] Mitsunori Miki, Tomoyuki Hiroyasu, and Jun’ya Wako. Adaptive Temperature Schedule Determined by Genetic Algorithm for Parallel Simulated Annealing. In *The 2003 Congress on Evolutionary Computation, CEC’03*, 2003. In proceedings [498]. Online available at http://mikilab.doshisha.ac.jp/dia/research/person/wako/society/cec2003/cec03_wako.pdf [accessed 2007-09-15].
 - [1332] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C++. Example Book. The Art of Scientific Computing*, chapter 10. Cambridge University Press, second edition, February 7, 2002. ISBN

- 0521750342, 978-0521750349. Partly online available at <http://books.google.de/books?id=gwijz-OyIYEC> [accessed 2008-06-14].
- [1333] Donald M. Olsson and Lloyd S. Nelson. The Nelder-Mead Simplex Procedure for Function Minimization. *Technometrics*, 17(1):45–51, February 1975.
- [1334] Jeffrey C. Lagarias, James A. Reeds, Margaret H. Wright, and Paul E. Wright. Convergence Properties of the Nelder-Mead Simplex Method in Low Dimensions. *SIAM Journal on Optimization (SIOPT)*, 9(1):112–147, 1998. ISSN 1052-6234 (print) / 1095-7189 (electronic). Online available at http://www.aoe.vt.edu/~cliff/aoe5244/nelder_mead_2.pdf [accessed 2008-06-14].
- [1335] Margaret H. Wright from the AT&T Bell Laboratories, Murray Hill, New Jersey, 07974, USA. Direct search methods: Once scorned, now respectable. In D. F. Griffiths and G. A. Watson, editors, *Numerical Analysis – Proceedings of the 16th Dundee Biennial Conference in Numerical Analysis*, Pitman Research Notes in Mathematics, pages 191–208, June 27–30, 1995, University of Dundee, Scotland, UK. Addison Wesley Longman Limited / Chapman & Hall/CRC, Harlow, UK. ISBN 0582276330, 978-0582276338. Online available at <http://citeseer.ist.psu.edu/155516.html> [accessed 2008-06-14].
- [1336] Michael W. Trosset from the Department of Computational, Applied Mathematics, Rice University, Houston, TX, USA. I Know It When I See It: Toward a Definition of Direct Search Methods. *SIAG/OPT Views and News: A Forum for the SIAM Activity Group on Optimization*, 9:7–10, Fall 1997. Online available at <http://citeseer.ist.psu.edu/135600.html> and <http://www.math.wm.edu/~trosset/Research/Opt/dsm3.ps.Z> [accessed 2008-06-14], online available at <http://www.mat.uc.pt/siagopt/vn9.pdf> [accessed 2008-07-02].
- [1337] K. I. M. McKinnon. Convergence of the Nelder-Mead simplex method to a nonstationary point. *SIAM Journal on Optimization*, 9(1):148–158, 1999. ISSN 1052-6234 (print) / 1095-7189 (electronic).
- [1338] Virginia Joanne Torczon. *Multi-Directional Search: A Direct Search Algorithm for Parallel Machines*. PhD thesis, Rice University, Houston, TX, USA, May 1989. Online available at <http://citeseer.ist.psu.edu/176729.html> [accessed 2008-06-14]. Supervisors: John E. Dennis, Jr., Andrew E. Boyd, Kenneth W. Kennedy, Jr., and Richard A. Tapia.
- [1339] Frederick H. Walters, Stephen L. Morgan, Lloyd R. Parker, Jr., and Stanley N. Deming. *Sequential Simplex Optimization: A Technique for Improving Quality and Productivity in Research, Development, and Manufacturing*. Chemometrics Series. CRC Press LLC, 2000 Corporate Blvd., N.W., Boca Raton, Florida 33431, USA, August 26, 1991. ISBN 0849358949, 978-0849358944. Licensed electronical reprint by MultiSimplex AB, Sternbergsg. 9, S-371 32 Karlskrona, Sweden. Online available at http://www.chem.sc.edu/faculty/morgan/pubs/WaltersParkerMorganDeming_SequentialSimplexOptimization.pdf [accessed 2008-06-14].
- [1340] Helio Joséé Correa Barbosa, Fernanda M. P. Raupp, Carlile Campos Lavor, H. Lima, and Nelson Maculan. A Hybrid Genetic Algorithm for Finding Stable Conformations of Small Molecules. In Felipe M. G. França and Carlos H. C. Ribeiro, editors, *Proceedings of the VI Brazilian Symposium on Neural Networks (SBRN 2000)*, pages 90–94, November 22–25, 2000, Rio de Janeiro, RJ, Brazil. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-0856-1. doi:10.1109/SBRN.2000.889719.
- [1341] Edmon Hok-Man Chan. Design and Implementation of a High Speed Cable-Based Planar Parallel Manipulator. Master’s thesis, University of Waterloo, Ontario, Canada, 2005. Advisor: Dr. Amir Khajepour. Online available at <http://etd.uwaterloo.ca/etd/hmechan2005.pdf> [accessed 2008-06-14].
- [1342] Simonetta Paloscia, Giovanni Macelloni, Paolo Pampaloni, and Emanuele Santi from the CNR-IFAC, via Panciatichi 64, 50127 Firenze, Italy. Retrieval of Soil

- Moisture Data at Global Scales with AMSR-E. In *XXVIIIth URSI General Assembly*, October 23–29, 2005, New Delhi, India. URSI Secretariat, c/o INTEC, Ghent University, Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium. Online available at [http://ursiweb.intec.ugent.be/Proceedings/ProcGA05/pdf/F06.2\(0590\).pdf](http://ursiweb.intec.ugent.be/Proceedings/ProcGA05/pdf/F06.2(0590).pdf) [accessed 2008-06-14].
- [1343] M. Musil, M. J. Wilmut, and N. R. Chapman from the School of Earth & Ocean Science, Victoria University, BC, Canada. A hybrid simplex genetic algorithm for estimating geoacoustic parameters using matched-field inversion. *IEEE Journal of Oceanic Engineering*, 24(3):358–369, July 1999. ISSN 0364-9059. doi:10.1109/48.775297. CODEN: IJOEDY, INSPEC Accession Number: 6322331.
- [1344] John Yen, James C. Liao, David Randolph, and Bogju Lee from the Department of Computer Science, Texas A&M University, College Station, TX, USA. A Hybrid Approach to Modeling Metabolic Systems Using Genetic Algorithms and the Simplex Method. In *CAIA '95: Proceedings of the 11th Conference on Artificial Intelligence for Applications*, pages 277–285, 1995. IEEE Computer Society, Washington, DC, USA. ISBN 0-8186-7070-3. doi:10.1109/CAIA.1995.378811. Online available at <http://citeseer.ist.psu.edu/33461.html> [accessed 2008-07-23].
- [1345] Suresh Chandra Satapathy, Jvr Murthy, P. V. G. D. Prasada Reddy, Venkatesh Katari, Satish Malireddi, and V. N. K. Srujan Kollisetty. An Efficient Hybrid Algorithm for Data Clustering Using Improved Genetic Algorithm and Nelder Mead Simplex Search. In *ICCIMA '07: Proceedings of the International Conference on Computational Intelligence and Multimedia Applications (ICCIMA 2007)*, pages 498–510, December 13–15, 2007, Sivakasi, Tamil Nadu, India. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-3050-8. doi:10.1109/ICCIMA.2007.68. INSPEC Accession Number: 9895376.
- [1346] Mordecai Avriel. *Nonlinear Programming: Analysis and Methods*. Dover Publishing Inc., September 16, 2003. ISBN 0-486-43227-0, 978-0486432274. Partly online available at <http://books.google.de/books?id=byF4Xb1QbvMC> [accessed 2008-06-14].
- [1347] Tetsuyuki Takahama and Setsuko Sakai from the Hiroshima City/Shudo University, Japan. Constrained optimization by applying the α constrained method to the nonlinear simplex method with mutations. *IEEE Transactions on Evolutionary Computation*, 9(5):437–451, October 3, 2005. ISSN 1089-778X. doi:10.1109/TEVC.2005.850256. INSPEC Accession Number: 8608336. Online available at <http://www.chi.its.hiroshima-cu.ac.jp/~takahama/eng/papers/aSimplex-TEC2005.pdf> [accessed 2008-07-23].
- [1348] Jürgen Jakumeit, Thomas Barth, Julian Reichwald, Manfred Grauer, Frank Thilo, Thomas Friese, Matthew Smith, and Bernd Freisleben. A Grid-based Parallel Optimization Algorithm Applied to a Problem in Metal Casting Industry. In *Proceedings of the Second International Conference on Bioinspired Optimization Methods and their Application, BIOMA 2006*, 2006. In proceedings [491].
- [1349] Thomas Barth, Bernd Freisleben, Manfred Grauer, and Frank Thilo from the Department of Inf. Syst., University of Siegen, Siegen, Germany. Distributed solution of optimal hybrid control problems on networks of workstations. In *Proceedings of the IEEE International Conference on Cluster Computing*, pages 162–169, November 28–December 1, 2000, Chemnitz, Germany. ISBN 0-7695-0896-0. doi:10.1109/CLUSTER.2000.889027.
- [1350] Fang Wang and Yuhui Qiu from the Intelligent Software and Software Engineering Laboratory, South-west China Normal University, Chongqing, China, 400715. Empirical Study of Hybrid Particle Swarm Optimizers with the Simplex Method Operator. In *ISDA '05: Proceedings of the 5th International Conference on Intelligent Systems Design and Applications*, pages 308–313, September 8–10, 2005, Wroclaw, Poland. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2286-06. doi:10.1109/ISDA.2005.44.

- [1351] Fang Wang, Yuhui Qiu, and Yun Bai. A New Hybrid NM Method and Particle Swarm Algorithm for Multimodal Function Optimization. In A. Fazel Famili, Joost N. Kok, José María Peña, Arno Siebes, and A. J. Feelders, editors, *Advances in Intelligent Data Analysis VI, 6th International Symposium on Intelligent Data Analysis, IDA 2005*, volume 3646 of *Lecture Notes in Computer Science (LNCS)*, pages 497–508. Springer, September 8–10, 2005, Madrid, Spain. ISBN 3-540-28795-7. doi:10.1007/11552253_45.
- [1352] Fang Wang and Yuhui Qiu from the Intelligent Software and Software Engineering Laboratory, South-west China Normal University, Chongqing, China, 400715. Multimodal Function Optimizing by a New Hybrid Nonlinear Simplex Search and Particle Swarm Algorithm. In João Gama, Rui Camacho, Pavel Brazdil, Alípio Jorge, and Luís Torgo, editors, *Proceedings of ECML 2005, the 16th European Conference on Machine Learning*, volume 3720 of *Lecture Notes in Artificial Intelligence (LNAI)*, subseries of *Lecture Notes in Computer Science (LNCS)*, pages 759–766. Springer, October 3–7 2005, Porto, Portugal. ISBN 3-540-29243-8. doi:10.1007/11564096_78.
- [1353] Fang Wang, Yuhui Qiu, and Naiqin Feng. Multi-model Function Optimization by a New Hybrid Nonlinear Simplex Search and Particle Swarm Algorithm. In *First International Conference on Advances in Natural Computation, Part III*, pages 562–565, 2005. doi:10.1007/11539902_68. In proceedings [319].
- [1354] Yuhui Qiu and Fang Wang from the Intelligent Software and Software Engineering Laboratory, South-west China Normal University, Chongqing, China, 400715. Improving Particle Swarm Optimizer Using the Nonlinear Simplex Method at Late Stage. In Richard T. Hurley and Wenying Feng, editors, *Proceedings of the ISCA, the 14th International Conference on Intelligent and Adaptive Systems and Software Engineering*, pages 25–30. ISCA, July 20–22, 2005, Novotel Toronto Centre, Toronto, Canada. ISBN 1-880843-55-2.
- [1355] Shu-Kai S. Fan, Yun-Chia Liang, and Erwie Zahara. A genetic algorithm and a particle swarm optimizer hybridized with Nelder-Mead simplex search. *Computers and Industrial Engineering*, 50(4):401–425, August 2006. ISSN 0360-8352. doi:10.1016/j.cie.2005.01.022. Special issue: Sustainability and globalization: Selected papers from the 32nd ICC&IE.
- [1356] Shu-Kai S. Fan and Erwie Zaharaa from the Department of Industrial Engineering and Management, Yuan Ze University, 135 Yuan-Tung Road, Chung-Li, Taoyuan County 320, Taiwan, ROC. A hybrid simplex search and particle swarm optimization for unconstrained optimization. *European Journal of Operational Research*, 181(2): 527–548, September 1 2007. doi:10.1016/j.ejor.2006.06.034. Continuous Optimization. Online available at <http://dx.doi.org/10.1016/j.ejor.2006.06.034> [accessed 2008-06-14].
- [1357] Changtong Luo and Bo Yu. Low Dimensional Simplex Evolution – A Hybrid Heuristic for Global Optimization. In *SNPD 2007: Eighth ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing*, volume 2, pages 470–474, July 30–August 1, 2007, Qingdao, Shandong, China. ISBN 978-0-7695-2909-7. doi:10.1109/SNPD.2007.58. INSPEC Accession Number: 9873104.
- [1358] Changtong Luo from the Department of Computational Science and Engineering, Graduate School of Engineering, Nagoya University, Furo-cho, Chikusa-ku, Nagoya 464-8603, Japan. *Low dimensional simplex evolution algorithms and their applications*. PhD thesis, Jilin University, 2007. in Chinese.
- [1359] Jean-Michel Renders and Hugues Bersini from the Lab. d’Autom., Univ. Libre de Bruxelles, Belgium. Hybridizing genetic algorithms with hill-climbing methods for global optimization: two possible ways. In *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence*,

- pages 312–317, 1994. doi:10.1109/ICEC.1994.349948. INSPEC Accession Number: 4818657. In proceedings [507].
- [1360] John Yen, David Randolph, Bogju Lee, and James C. Liao from the Department of Computer Science, Texas A&M University, College Station, TX, USA. A hybrid genetic algorithm for the identification of metabolic models. In *TAI'95: Proceedings of the Seventh International Conference on Tools with Artificial Intelligence*, pages 4–7, November 5–8, 1995, Herndon, VA, USA. IEEE Computer Society, Washington, DC, USA. ISBN 0-8186-7312-5. doi:10.1109/TAI.1995.479371. INSPEC Accession Number: 5162198.
- [1361] John Yen and Bogju Lee from the Department of Computer Science, Texas A&M University, College Station, TX, USA. A simplex genetic algorithm hybrid. In *IEEE International Conference on Evolutionary Computation*, pages 175–180, 1997. IEEE Computer Society, Washington, DC, USA. doi:10.1109/ICEC.1997.592291. INSPEC Accession Number: 5573047. In proceedings [504].
- [1362] Guoli Zhang, Hai Yan Lu, Gengyin Li, and Hong Xie from the North China Electr. Power Univ. A New Hybrid Real-Coded Genetic Algorithm and Application in Dynamic Economic Dispatch. In *WCICA 2006. The Sixth World Congress on Intelligent Control and Automation*, volume 1, pages 3627–3632, June 21–23, 2006, Dalian, China. ISBN 1-4244-0332-4. doi:10.1109/WCICA.2006.1713046. INSPEC Accession Number: 9187799.
- [1363] Guoli Zhang and Haiyan Lu. Hybrid Real-Coded Genetic Algorithm with Quasi-Simplex Technique. *IJCSNS International Journal of Computer Science and Network Security*, 6(10):246–255, October 2006. ISSN 1738-7906. Online available at http://paper.ijcsns.org/07_book/200610/200610B15.pdf and [accessed 2008-07-29].
- [1364] Aesop Rock. Garbage. In *Float*, Aesop Rock, Omega One, and Blockhead, producers. Mush Records, November 6, 2000.
- [1365] Fred Glover. Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, 13(5):533–549, 1986. ISSN 0305-0548. doi:10.1016/0305-0548(86)90048-1. Online available at [http://dx.doi.org/10.1016/0305-0548\(86\)90048-1](http://dx.doi.org/10.1016/0305-0548(86)90048-1) [accessed 2008-03-27].
- [1366] P. Hansen. The steepest ascent mildest descent heuristic for combinatorial programming. In *Talk presented at the Congress on Numerical Methods in Combinatorial Optimization*. Academic Press, 1986, Capri, Italy.
- [1367] Fred Glover. Tabu Search – Part I. *Operations Research Society of America (ORSA) Journal on Computing*, 1(3):190–206, Summer 1989. Online available at <http://leeds-faculty.colorado.edu/glover/TS%20-%20Part%20I-ORSA.pdf> [accessed 2008-03-27].
- [1368] Fred Glover. Tabu Search – Part II. *Operations Research Society of America (ORSA) Journal on Computing*, 2(1):190–206, Winter 1990. Online available at <http://leeds-faculty.colorado.edu/glover/TS%20-%20Part%20II-ORSA.pdf> [accessed 2008-03-27].
- [1369] Dominique de Werra and Alain Hertz. Tabu Search Techniques: A tutorial and an application to neural networks. *OR Spectrum*, 11:131–141, 1989. ISSN 0171-6468 (Print) 1436-6304 (Online). Online available at <http://www.springerlink.de/content/x25k97k0qx237553/fulltext.pdf> [accessed 2008-03-27].
- [1370] Alain Hertz, Eric Taillard, and Dominique de Werra. A Tutorial on Tabu Search. In *Proceedings of Giornate di Lavoro AIRO'95 (Enterprise Systems: Management of Technological and Organizational Changes)*, pages 13–24, 1995, Italy. Online available at <http://citeseer.ist.psu.edu/73006.html> [accessed 2007-09-15].
- [1371] R. Battiti and G. Tecchiolli. The reactive tabu search. *ORSA Journal on Computing*, 6(2):126–140, 1994. Online available at <http://citeseer.ist.psu.edu/141556.html> and <http://rtm.science.unitn.it/~battiti/archive/reactive-tabu-search.ps.gz> [accessed 2007-09-15].

- [1372] Djurdje Cvijović and Jacek Klinowski from the Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, UK. Taboo search: an approach to the multiple minima problem. *Science*, 267(5198):664–666, February 3, 1995. ISSN 0036-8075. doi:10.1126/science.267.5198.664.
- [1373] Buyang Cao and Fred Glover. Tabu Search and Ejection Chains-Application to a Node Weighted Version of the Cardinality-Constrained TSP. *Management Science*, 43(7):908–921, July 1997. ISSN 0025-1909.
- [1374] Johannes J. Schneider and Scott Kirkpatrick, editors. *Tabu Search Applied to TSP*, chapter part II (Applications), chapter 1, pages 441–447. Scientific Computation. Springer Berlin Heidelberg, 2006. ISBN 978-3-540-34559-6 (Print) 978-3-540-34560-2 (Online). doi:10.1007/978-3-540-34560-2_42.
- [1375] John F. McLoughlin III and Walter Cedeño. The enhanced evolutionary tabu search and its application to the quadratic assignment problem. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 975–982, 2005. doi:10.1145/1068009.1068175. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068175> [accessed 2007-08-25].
- [1376] Fred Glover and M. Laguna. Tabu Search. In *Modern Heuristic Techniques for Combinatorial Problems*. John Wiley & Sons, 1993. ISBN 978-0470220795. In collection [406]. Online available at <http://citeseer.ist.psu.edu/glover97tabu.html> [accessed 2007-09-16].
- [1377] Fred Glover and M. Laguna. Tabu search. In *Handbook of Combinatorial Optimization*, volume 3, pages 621–757. Kluwer Academic Publishers, Springer Netherlands, 1998. In collection [405].
- [1378] Gwo-Jen Hwang, Peng-Yeng Yin, and Shu-Heng Yeh. A tabu search approach to generating test sheets for multiple assessment criteria. *IEEE Transactions on Education*, 49(1):88–97, February 2006. doi:10.1109/TE.2002.858405.
- [1379] Paolo Toth and Daniele Vigo. The Granular Tabu Search and Its Application to the Vehicle-Routing Problem. *INFORMS Journal on Computing*, 15(4):333–346, 2003. ISSN 1526-5528. doi:10.1287/ijoc.15.4.333.24890.
- [1380] Randall S. Sexton, Bahram Alidaee, Robert E. Dorsey, and John D. Johnson. Global optimization for artificial neural networks: A tabu search application. *European Journal of Operational Research*, 106(2–3):570–584, April 1998. doi:10.1016/S0377-2217(97)00292-0. Online available at [http://dx.doi.org/10.1016/S0377-2217\(97\)00292-0](http://dx.doi.org/10.1016/S0377-2217(97)00292-0) [accessed 2007-08-25].
- [1381] M. Szachniuk, L. Popenda, Z. Gdaniec, R.W. Adamiak, and J. Błażewicz from the Institute of Bioorganic Chemistry, Polish Academy of Sciences, Noskowskiego 12/14, 61-704 Poznan, Poland. NMR Analysis of RNA Bulged Structures: Tabu Search Application in NOE Signal Assignment. In *Proceedings of the 2005 IEEE Symposium on Computational Intelligence in Bioinformatics and Computational Biology CIBCB'05*, pages 172–178. IEEE, November 2005, San Diego, USA. Online available at <http://www.man.poznan.pl/~lpopenda/HomePage/pub/CIBCB2005172.pdf> [accessed 2007-08-25].
- [1382] Andreas Fink and S. Voß. Generic Application of Tabu Search Methods to Manufacturing Problems. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC'98)*, volume 3, pages 2385–2390. IEEE, October 1998, Piscataway. doi:10.1109/ICSMC.1998.725013.
- [1383] Samuel Pierre and A. Elgibaoui. Improving communication network topologies using tabu search. In *LCN: Proceedings of the 22nd IEEE Conference on Local Computer Networks (LCN '97)*, pages 44–53. IEEE Computer Society, November 2–5 1997, Minneapolis, Minnesota, USA. ISBN 0-8186-8141-1.
- [1384] Samuel Pierre and Fabien Houéto from the Department of Computer Engineering, Mobile Computing, and Networking Research (LARIM), École Polytechnique de Montréal, C.P. 6079, Succ. Centre-Ville, Montréal, Québec, Canada H3C 3A7. A tabu search approach for assigning cells to switches in cellular mobile networks. *Computer*

- Communications*, 25(5):464–477, March 2002. doi:10.1016/S0140-3664(01)00371-1. Online available at [http://dx.doi.org/10.1016/S0140-3664\(01\)00371-1](http://dx.doi.org/10.1016/S0140-3664(01)00371-1) [accessed 2008-08-01].
- [1385] Samuel Pierre and Fabien Houéto from the Department of Computer Engineering, Mobile Computing, and Networking Research (LARIM), École Polytechnique de Montréal, C.P. 6079, Succ. Centre-Ville, Montréal, Québec, Canada H3C 3A7. Assigning cells to switches in cellular mobile networks using taboo search. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 32(3):351–356, June 2002. doi:10.1109/TSMCB.2002.999810.
- [1386] Marc St-Hilaire, Steven Chamberland, and Samuel Pierre. A Tabu Search Heuristic for the Global Planning of UMTS Networks. In *Proceedings of the IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob'2006)*, pages 148–151, June 19–21, 2006, Montréal, Québec, Canada. ISBN 1-4244-0494-0. doi:10.1109/WIMOB.2006.1696339. INSPEC Accession Number: 9132969.
- [1387] Jean-Louis Deneubourg, Jacques M. Pasteels, and J. C. Verhaeghe. Probabilistic behaviour in ants: a strategy of errors? *Journal of Theoretical Biology*, 105(2):259–271, 1983. doi:10.1016/S0022-5193(83)80007-1. Online available at [http://dx.doi.org/10.1016/S0022-5193\(83\)80007-1](http://dx.doi.org/10.1016/S0022-5193(83)80007-1) [accessed 2008-03-27].
- [1388] Jean-Louis Deneubourg and Simon Goss. Collective patterns and decision-making. *Ethology, Ecology & Evolution*, 1(4):295–311, December 1989.
- [1389] Simon Goss, R. Beckers, Jean-Louis Deneubourg, S. Aron, and Jacques M. Pasteels. How trail laying and trail following can solve foraging problems for ant colonies. In R. N. Hughes, editor, *Behavioural Mechanisms of Food Selection*, NATO ASI Series, G 20, pages 661–678. Springer-Verlag, Berlin, April 1990. ISBN 0387517626, 978-0387517629.
- [1390] T. R. Stickland, C. M. N. Tofts, and N. R. Franks from the School of Biological Sciences, University of Bath, Claverton Down, Bath, Avon, UK. A path choice algorithm for ants. *Naturwissenschaften*, 79(12):567–572, December 1992. ISSN 0028-1042 (Print) 1432-1904 (Online). doi:10.1007/BF01131415. Online available at <http://www.springerlink.com/content/v2t8328844843156/fulltext.pdf> [accessed 2008-06-12].
- [1391] Marco Dorigo, Vittorio Maniezzo, and Alberto Coloni. The Ant System: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics Part B: Cybernetics*, 26(1):29–41, 1996. Online available at <ftp://iridia.ulb.ac.be/pub/mdorigo/journals/IJ.10-SMC96.pdf> and <http://citeseer.ist.psu.edu/dorigo96ant.html> [accessed 2007-08-05].
- [1392] Marco Dorigo and Christian Blum. Ant colony optimization theory: a survey. *Theoretical Computer Science*, 344(2-3):243–278, 2005. ISSN 0304-3975. Online available at <http://code.ulb.ac.be/dbfiles/DorBlu2005tcs.pdf> [accessed 2007-08-05].
- [1393] Marco Dorigo, Gianni Di Caro, and L. Gambardella. Ant Algorithms for Discrete Optimization. Technical Report IRIDIA/98-10, Université Libre de Bruxelles, Belgium, 1998. Online available at <http://citeseer.ist.psu.edu/dorigo98ant.html> [accessed 2007-08-05]. See [1946].
- [1394] Bernard Manderick and Frans Moyson. The collective behaviour of Ants: an Example of Self-Organization in Massive Parallelism. In *Proceedings of AAAI Spring Symposium on Parallel Models of Intelligence*, 1988, Stanford, California.
- [1395] Vittorio Maniezzo, Luca Maria Gambardella, and Fabio de Luigi. Ant Colony Optimization. In *New Optimization Techniques in Engineering*, chapter 5, pages 101–117. Springer-Verlag, 2004. In collection [404]. Online available at <http://www.idsia.ch/~luca/aco2004.pdf> and <http://citeseer.ist.psu.edu/maniezzo04ant.html> [accessed 2007-09-13].

- [1396] Pierre-Paul Grassé. La reconstruction du nid et les coordinations inter-individuelles chez *bellicositermes natalensis* et *cubitermes* sp. la théorie de la stigmergie: Essai d'interprétation des termites constructeurs. *Insectes Sociaux*, 6:41–48, 1959, Paris.
- [1397] Ruud Schoonderwoerd, Owen E. Holland, Janet L. Bruten, and Leon J. M. Rothkrantz. Ant-Based Load Balancing in Telecommunications Networks. *Adaptive Behavior*, 5(2):169–207, 1996. Online available at <http://citeseer.ist.psu.edu/schoonderwoerd96antbased.html> and <http://www.ifi.unizh.ch/ailab/teaching/FG06/> [accessed 2007-08-19].
- [1398] Edward Osborne Wilson. *Sociobiology: The New Synthesis*. Belknap Press / Harvard University Press, Cambridge, Massachusetts, USA / Cumbreland, Rhode Island, USA, 1975. See also [1947].
- [1399] Guy Théraulaz and Eric Bonabeau. Coordination in Distributed Building. *Science*, 269(5224):686–688, August 4, 1995. doi:10.1126/science.269.5224.686.
- [1400] Justin Werfel and Radhika Nagpal. Extended Stigmergy in Collective Construction. *IEEE Intelligent Systems*, 21(2):20–28, March/April 2006. doi:10.1109/MIS.2006.25. Online available at <http://hebb.mit.edu/people/jkwerfel/ieeais06.pdf> [accessed 2008-06-12].
- [1401] Crystal Schuil, Matthew Valente, Justin Werfel, and Radhika Nagpal. Collective Construction Using Lego Robots. In *Annual National Conference on Artificial Intelligence (AAAI), Robot Exhibition*, 2006. In proceedings [279]. Online available at <http://www.eecs.harvard.edu/~rad/ssr/papers/aaai06-schuil.pdf> [accessed 2008-06-12].
- [1402] Justin Werfel, Yaneer Bar-Yam, Daniela Rus, and Radhika Nagpal. Distributed Construction by Mobile Robots with Enhanced Building Blocks. In *IEEE International Conference on Robotics and Automation (ICRA)*, May 15–19, 2006, Hilton in the Walt Disney World Resort hotel, Walt Disney World Resort, Orlando, Florida (Lake Buena Vista), USA. Online available at <http://www.eecs.harvard.edu/~rad/ssr/papers/icra06-werfel.pdf> [accessed 2008-06-12].
- [1403] Peter Korošec and Juri Šilc. Real-Parameter Optimization Using Stigmergy. In *Proceedings of the Second International Conference on Bioinspired Optimization Methods and their Application, BIOMA 2006*, pages 73–84, October 2006. Online available at <http://csd.ijs.si/silc/articles/BIOMA06DASA.pdf> [accessed 2007-08-05].
- [1404] Luca Maria Gambardella and Marco Dorigo. Solving Symmetric and Asymmetric TSPs by Ant Colonies. In *International Conference on Evolutionary Computation*, pages 622–627, 1996. In proceedings [505]. Online available at <http://citeseer.ist.psu.edu/gambardella96solving.html> and <http://www.idsia.ch/~luca/icec96-acsc.pdf> [accessed 2007-08-05].
- [1405] Marco Dorigo and Luca Maria Gambardella. Ant Colony System: A Cooperative Learning Approach to the Traveling Salesman Problem. *IEEE Transactions on Evolutionary Computation*, 1(1):53–66, April 1997. Online available at <http://citeseer.ist.psu.edu/153.html> and <http://www.idsia.ch/~luca/acsc-ec97.pdf> [accessed 2007-08-19].
- [1406] Luca Maria Gambardella, Andrea E. Rizzoli, Fabrizio Oliverio, Norman Casagrande, Alberto V. Donati, Roberto Montemanni, and Enzo Lucibello. Ant Colony Optimization for vehicle routing in advanced logistics systems. In *Proceedings of MAS 2003 International Workshop on Modelling and Applied Simulation*, pages 3–9, October 2003, Bergeggi, Italy. Online available at http://www.idsia.ch/~luca/MAS2003_18.pdf [accessed 2007-08-19].
- [1407] Luca Maria Gambardella, É. D. Taillard, and Marco Dorigo. Ant Colonies for the Quadratic Assignment Problem. *The Journal of the Operational Research Society*, 50(2):167–176, February 1999. doi:10.2307/3010565. Online available at <http://www.idsia.ch/~luca/tr-idsia-4-97.pdf> and <http://citeseer.ist.psu.edu/378986.html> [accessed 2007-08-19].

- [1408] Daniel Merkle, Martin Middendorf, and Hartmut Schneck. Ant Colony Optimization for Resource-Constrained Project Scheduling. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*, pages 893–900, 2000. In proceedings [570]. Online available at <http://citeseer.ist.psu.edu/merkle00ant.html> and <http://pacosy.informatik.uni-leipzig.de/pv/Personen/middendorf/Papers/> [accessed 2007-08-19].
- [1409] Ruud Schoonderwoerd. Collective intelligence for network control. Master's thesis, Delft University of Technology, Faculty of Technical Mathematics and Informatics, May 30, 1996. Online available at <http://staff.washington.edu/paymana/swarm/schoon96-thesis.pdf> [accessed 2008-06-12]. Graduation committee: Prof. dr. H. Koppelaar, Prof. dr. ir. E.J.H. Kerckhoffs, Drs. dr. L.J.M. Rothkrantz, and J.L. Bruten M.Sc.
- [1410] Guy Théraulaz and Eric Bonabeau. Swarm Smarts. *Scientific American*, pages 72–79, March 2000. Online available at <http://www.santafe.edu/~vince/press/SciAm-SwarmSmarts.pdf> [accessed 2008-06-12].
- [1411] K.M. Sim and W.H. Sun. Multiple Ant-Colony Optimization for Network Routing. In *First International Symposium on Cyber Worlds (CW'02)*, page 0277, 2002. IEEE Computer Society, Los Alamitos, CA, USA. ISBN 0-7695-1862-1. doi:10.1109/CW.2002.1180890.
- [1412] Agostino Forestiero, Carlo Mastroianni, and Giandomenico Spezzano from the Institute for High Performance Computing and Networks (ICAR-CNR), 87036 Rende (CS) Italy. Antares: an Ant-Inspired P2P Information System for a Self-Structured Grid. In *Proceedings of BIONETICS 2007*, 2007. In proceedings [1936]. Online available at <http://grid.deis.unical.it/papers/pdf/ForestieroMastroianniSpezzano-Antares.pdf> [accessed 2008-06-13].
- [1413] Suiliong Liang, A. Nur Zincir-Heywood, and Malcom I. Heywood from the Faculty of Computer Science, Dalhousie University, Halifax, NS, Canada. The effect of routing under local information using a social insect metaphor. In *CEC'02: Proceedings of the Congress on Evolutionary Computation*, pages 1438–1443, 2002. IEEE Computer Society, Washington, DC, USA. In proceedings [499]. Online available at <http://users.cs.dal.ca/~mheywood/X-files/Publications/01004454.pdf> [accessed 2008-07-26].
- [1414] Gianni Di Caro from the IRIDIA, Université Libre de Bruxelles, 50, av. F. Roosevelt, CP 194/6, 1050, Brussels, Belgium. *Ant Colony Optimization and its application to adaptive routing in telecommunication networks*. PhD thesis, Applied Sciences, Polytechnic School, Université Libre de Bruxelles, Brussels, Belgium, September 2004. Advisor: Marco Dorigo. Partly online available at <http://www.idsia.ch/~gianni/Papers/thesis-abstract.pdf> [accessed 2008-07-26] and <http://www.idsia.ch/~gianni/Papers/routing-chapters.pdf> [accessed 2008-07-26].
- [1415] Suiliong Liang, A. Nur Zincir-Heywood, and Malcolm I. Heywood from the NIMS Group, Faculty of Computer Science, Dalhousie University, 6050 University Avenue, Halifax, NS, Canada B3H 1W5. Adding more intelligence to the network routing problem: AntNet and Ga-agents. *Applied Soft Computing*, 6(3):244–257, March 2006. ISSN 1568-4946. doi:10.1016/j.asoc.2005.01.005.
- [1416] Gianni Di Caro and Marco Dorigo from the IRIDIA, Université Libre de Bruxelles 50, av. F. Roosevelt, CP 194/6 1050, Brussels, Belgium. Two ant colony algorithms for best-effort routing in datagram networks. In *Proceedings of the Tenth IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS'98)*, pages 541–546. IASTED/ACTA Press, October 1–4, 1998, Las Vegas, USA. Online available at <http://citeseer.ist.psu.edu/dicaro98two.html> and <http://www.idsia.ch/~gianni/Papers/PDCS98.ps.gz> [accessed 2008-07-26].
- [1417] Gianni Di Caro and Marco Dorigo from the IRIDIA, Université Libre de Bruxelles, 50, av. F. Roosevelt, CP 194/6, 1050, Brussels, Belgium. AntNet: Distributed Stig-

- mergetic Control for Communications Networks. *Journal of Artificial Intelligence Research (JAIR)*, 9:317–365, 1998. ISSN 11076-9757. Submitted 5/98; published 12/98. Online available at <http://citeseer.ist.psu.edu/dicaro98antnet.html> and <http://www.jair.org/media/544/live-544-1748-jair.pdf> [accessed 2008-07-26].
- [1418] Russel C. Eberhart and James Kennedy. A New Optimizer Using Particle Swarm Theory. In *Proceedings of the Sixth International Symposium on Micro Machine and Human Science MHS'95*, pages 39–43. IEEE Press, October 1995. ISBN 0-7803-2676-8.
- [1419] James Kennedy and Russel C. Eberhart. Particle swarm optimization. In *Proceedings of IEEE International Conference on Neural Networks, 1995*, volume 4, pages 1942–1948, November 27–December 1, 1995, Perth, WA, Australia. ISBN 0-7803-2768-3. Online available at <http://www.engr.iupui.edu/~shi/Coference/psopap4.html> [accessed 2007-08-21].
- [1420] Gerhard Venter and Jaroslaw Sobieszczanski-Sobieski. Particle Swarm Optimization. *AIAA Journal*, 41(8):1583–1589, 2003. AIAA 2002-1235. 43rd AIAA/ASME/ASCE/AHS/ASC, Structures, Structural Dynamics, and Materials Conference, April 22-25, 2002, Denver, Colorado. Online available at <http://citeseer.ist.psu.edu/venter02particle.html> [accessed 2007-08-20].
- [1421] Tao Cai, Feng Pan, and Jie Chen. Adaptive particle swarm optimization algorithm. In *Proceedings of Fifth World Congress on Intelligent Control and Automation, 2004. WCICA 2004*, volume 3, pages 2245–2247, June 15–19, 2004. ISBN 0-7803-8273-0. doi:10.1109/WCICA.2004.134198.
- [1422] Yuelin Gao and Yuhong Duan. An Adaptive Particle Swarm Optimization Algorithm with New Random Inertia Weight. In *Advanced Intelligent Computing Theories and Applications. With Aspects of Contemporary Intelligent Computing Techniques*, volume 2, pages 342–350, 2007. doi:10.1007/978-3-540-74282-1. In proceedings [1498].
- [1423] Yuelin Gao and Zihui Ren. Adaptive Particle Swarm Optimization Algorithm With Genetic Mutation Operation. In *Proceedings of the Third International Conference on Natural Computation (ICNC 2007)*, volume 2, pages 211–215, 2007. doi:10.1109/ICNC.2007.161. In proceedings [311].
- [1424] Russel C. Eberhart and Yuhui Shi. Comparison between Genetic Algorithms and Particle Swarm Optimization. In *Proceedings of the 7th International Conference on Evolutionary Programming*, pages 611–616, 1998. In proceedings [1186].
- [1425] Peter John Angeline. Evolutionary Optimization Versus Particle Swarm Optimization: Philosophy and Performance Differences. In *EP'98: Proceedings of the 7th International Conference on Evolutionary Programming VII*, pages 601–610, 1998. In proceedings [1186].
- [1426] Michael Meissner, Michael Schmuker, and Gisbert Schneider from the Johann Wolfgang Goethe-Universität, Institut für Organische Chemie und Chemische Biologie, Siesmayerstraße 70, D-60323 Frankfurt, Germany. Optimized Particle Swarm Optimization (OPSO) and its application to artificial neural network training. *BMC Bioinformatics*, 7(125), 2006. doi:10.1186/1471-2105-7-125. Online available at <http://www.biomedcentral.com/1471-2105/7/125> [accessed 2007-08-21].
- [1427] Thomas Kiel Rasmussen and Thiemo Krink. Improved Hidden Markov Model training for multiple sequence alignment by a particle swarm optimization-evolutionary algorithm hybrid. *BioSystems*, 72(1–2):5–17, November 2003. Online available at [http://dx.doi.org/10.1016/S0303-2647\(03\)00131-X](http://dx.doi.org/10.1016/S0303-2647(03)00131-X) [accessed 2007-08-21].
- [1428] K. E. Parsopoulos and M. N. Vrahatis. Recent approaches to global optimization problems through Particle Swarm Optimization. *Natural Computing*, 1(2–3):235–306, June 2002. Online available at <http://citeseer.ist.psu.edu/parsopoulos02recent.html> [accessed 2007-08-21].
- [1429] Ge Wu, Volkert Hansen, E. Kreysa, and H.-P. Gemünd. Optimierung von FSS-Bandpassfiltern mit Hilfe der Schwarmintelligenz (Particle Swarm Optimization).

- Advances in Radio Science*, 4:65–71, September 2006. Online available at <http://www.adv-radio-sci.net/4/65/2006/ars-4-65-2006.pdf> [accessed 2007-08-21].
- [1430] Alexandre M. Baltar and Darrell G. Fontane. A generalized multiobjective particle swarm optimization solver for spreadsheet models: application to water quality. In *Proceedings of AGU Hydrology Days 2006*, pages 1–12, March 20–22, 2006, Fort Collins, Colorado. Online available at <http://www.lania.mx/~ccoello/EM00/baltar06.pdf.gz> and http://hydrologydays.colostate.edu/Proceedings_2006.htm [accessed 2007-08-21].
- [1431] Walter Cedeño and Dimitris K. Agrafiotis. Using particle swarms for the development of QSAR models based on K-nearest neighbor and kernel regression. *Journal of Computer-Aided Molecular Design*, 17(2–4):255–263, February 2003. ISSN 0920-654X (Print) 1573-4951 (Online). doi:10.1023/A:1025338411016. Online available at <http://www.springerlink.com/content/j523757110202636/> and <http://www.dimitris-agrafiotis.com/> [accessed 2007-08-21].
- [1432] Qi Shen, Jian-Hui Jiang, Chen-Xu Jiao, Shuang-Yan Huan, Guo li Shen, , and Ru-Qin Yu. Optimized Partition of Minimum Spanning Tree for Piecewise Modeling by Particle Swarm Algorithm. QSAR Studies of Antagonism of Angiotensin II Antagonists. *Journal of Chemical Information and Modeling / J. Chem. Inf. Comput. Sci.*, 44(6): 2027–2031, November 2004. doi:10.1021/ci034292+ S0095-2338(03)04292-6. Online available at <http://dx.doi.org/10.1021/ci034292+> [accessed 2007-08-21].
- [1433] Yuhui Shi and Marco Dorigo, editors. *Proceedings of the 2007 IEEE Swarm Intelligence Symposium. SIS'07*, IEEE International Symposia on Swarm Intelligence, April 1–5, 2007, Hilton Hawaiian Village Beach Resort & Spa, Hawaii, USA. Institute of Electrical & Electronics Engineer (IEEE). <http://www.computelligence.org/sis/2007/> [accessed 2007-08-26].
- [1434] *Proceedings of the 2006 IEEE Swarm Intelligence Symposium. SIS'06*, IEEE International Symposia on Swarm Intelligence, May 12–14, 2006, Hilton Indianapolis Hotel, Indianapolis, Indiana, USA. Institute of Electrical & Electronics Engineer (IEEE). <http://www.computelligence.org/sis/2006/> [accessed 2007-08-26].
- [1435] *Proceedings of the 2005 IEEE Swarm Intelligence Symposium. SIS2005*, IEEE International Symposia on Swarm Intelligence, June 8–10, 2005, The Westin Pasadena, Pasadena, California, USA. Institute of Electrical & Electronics Engineer (IEEE). ISBN 978-0780389168.
- [1436] *Proceedings of the 2003 IEEE Swarm Intelligence Symposium. SIS'03*, IEEE International Symposia on Swarm Intelligence, April 24–26, 2003, University Place Conference Center, Indiana University Purdue University Indianapolis, Indianapolis, Indiana, USA. Institute of Electrical & Electronics Engineer (IEEE). ISBN 978-0780379145. doi:10.1109/SIS.2003.1202237. Cat. No.03EX706, <http://www.computelligence.org/sis/2003/> [accessed 2007-08-26].
- [1437] Nadia Nedjah and Luiza de Macedo Mourelle, editors. *Swarm Intelligent Systems*, volume 26 of *Studies in Computational Intelligence*. Springer, Berlin/Heidelberg/New York, July 28, 2006. ISBN 3-540-33868-3, 978-3-540-33868-0. Series editor: Janusz Kacprzyk.
- [1438] Pablo Moscato. On Evolution, Search, Optimization, Genetic Algorithms and Martial Arts: Towards Memetic Algorithms. Technical Report C3P 826, Caltech Concurrent Computation Program 158-79, California Institute of Technology, Pasadena, CA 91125, USA, Pasadena, CA, 1989. Online available at <http://www.densis.fee.unicamp.br/~moscato/papers/bigone.ps> [accessed 2007-08-18].
- [1439] Richard Dawkins. *The Selfish Gene*. Oxford University Press, first: 1976, second edition, October 1989. ISBN 0192860925, 978-0192860927, 0-192-86092-5. Partly online available at <http://books.google.de/books?id=WkH09HI7koEC> [accessed 2008-03-28].

- [1440] Jason Digalakis and Konstantinos Margaritis. Performance Comparison of Memetic Algorithms. *Journal of Applied Mathematics and Computation*, 158:237–252, October 2004. Online available at <http://citeseer.ist.psu.edu/458892.html> and <http://www.complexity.org.au/ci/draft/draft/digala02/digala02s.pdf> [accessed 2007-09-12].
- [1441] Jason Digalakis and Konstantinos Margaritis. A parallel memetic algorithm for solving optimization problems. In *Proceedings of the 4th Metaheuristics International Conference*, pages 121–125, 2001. In proceedings [388]. Online available at <http://citeseer.ist.psu.edu/digalakis01parallel.html> and <http://www.it.uom.gr/people/digalakis/digiasfull.pdf> [accessed 2007-09-12].
- [1442] Maria J. Blesa, Pablo Moscato, and Fatos Xhafa. A Memetic Algorithm for the Minimum Weighted k -Cardinality Tree Subgraph Problem. In *Proceedings of the 4th Metaheuristics International Conference*, pages 85–90, 2001. In proceedings [388]. Online available at <http://citeseer.ist.psu.edu/458772.html> and <http://citeseer.ist.psu.edu/649471.html> [accessed 2007-09-12].
- [1443] Luciana Buriol, Paulo M. França, and Pablo Moscato. A New Memetic Algorithm for the Asymmetric Traveling Salesman Problem. *Journal of Heuristics*, 10(5):483–506, September 2004. ISSN 1381-1231 (Print) 1572-9397 (Online). doi:10.1023/B:HEUR.0000045321.59202.52. Online available at <http://www.springerlink.com/content/w617486q60mphg88/fulltext.pdf> and <http://citeseer.ist.psu.edu/544761.html> [accessed 2007-09-12].
- [1444] Natalio Krasnogor and Jim Smith. A tutorial for competent memetic algorithms: model, taxonomy, and design issues. *IEEE Transactions on Evolutionary Computation*, 9(5):474–488, October 2005. doi:10.1109/TEVC.2005.850260. Online available at <http://www.cs.nott.ac.uk/~nxk/PAPERS/IEEE-TEC-lastVersion.pdf> [accessed 2008-03-28].
- [1445] Peter Merz and Bernd Freisleben. A Comparison of Memetic Algorithms, Tabu Search, and Ant Colonies for the Quadratic Assignment Problem. In *Proceedings of the Congress on Evolutionary Computation*, volume 3, pages 2063–2070, 1999. In proceedings [502], Online available at <http://en.scientificcommons.org/204950> and <http://citeseer.ist.psu.edu/merz99comparison.html> [accessed 2008-03-28].
- [1446] William E. Hart, Natalio Krasnogor, and J. E. Smith, editors. *Recent Advances in Memetic Algorithms*. Studies in Fuzziness and Soft Computing. Springer, Berlin, Germany, September 2007. ISBN 3540229043, 978-3540229049. Partly online available at <http://books.google.de/books?id=LYf7YW4DmkUC> [accessed 2008-04-01].
- [1447] Xin Yao, F. Wang, K. Padmanabhan, and Sancho Salcedo-Sanz. Hybrid Evolutionary Approaches to Terminal Assignment in Communication Networks. In *Evolutionary Multiobjective Optimization – Theoretical Advances and Applications*. Springer, 2005. In collection [614].
- [1448] Alejandro Quintero and Samuel Pierre from the Department of Computer Engineering, Mobile Computing, and Networking Research (LARIM), École Polytechnique de Montréal, C.P. 6079, Succ. Centre-Ville, Montréal, Québec, Canada H3C 3A7. Sequential and multi-population memetic algorithms for assigning cells to switches in mobile networks. *Computer Networks*, 43(3):247–261, October 22, 2003. doi:10.1016/S1389-1286(03)00270-6. Online available at [http://dx.doi.org/10.1016/S1389-1286\(03\)00270-6](http://dx.doi.org/10.1016/S1389-1286(03)00270-6) [accessed 2008-08-01].
- [1449] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. MIT Electrical Engineering and Computer Science. The MIT Press and McGraw-Hill, second edition, August 2001. ISBN 0-2625-3196-8, first edition: 978-0262031417. First edition June 1990.
- [1450] Werner Dilger. *Einführung in die Künstliche Intelligenz*. Chemnitz University of Technology, Faculty of Computer Science, Chair of Artificial Intelligence (Künstliche Intelligenz), April 2006. Lecture notes for the lectures artificial intelligence. Online

- available at <http://www.tu-chemnitz.de/informatik/KI/skripte.php> [accessed 2007-08-06].
- [1451] Dawn Xiaodong Song, Adrian Perrig, and Doantam Phan from the University of California, Berkeley, USA. AGVI – Automatic Generation, Verification, and Implementation of Security Protocols. In *Computer Aided Verification – CAV’01: Proceedings of the 13th International Conference on Computer Aided Verification*, volume 2102/2001 of *Lecture Notes in Computer Science (LNCS)*, pages 241–245. Springer Berlin / Heidelberg, July 18–22, 2001, Paris, France. ISBN 3-540-42345-1, 978-3-540-42345-4. doi:10.1007/3-540-44585-4_21.
 - [1452] Adrian Perrig and Dawn Xiaodong Song from the Computer Science Department of Carnegie Mellon University and University of California, Berkeley, USA. A First Step towards the Automatic Generation of Security Protocols. In *In Proceedings of the Symposium on Network and Distributed Systems Security (NDSS)*. The Internet Society (ISOC), February 2–4, 2000, Catamaran Resort Hotel, San Diego, California, USA. ISBN 1-891562-07-X, 1-891562-08-8. Online available at <http://www.cs.berkeley.edu/~dawnsong/papers/apg.pdf> and <http://www.isoc.org/isoc/conferences/ndss/2000/proceedings/031.pdf> [accessed 2008-06-23].
 - [1453] Marc St-Hilaire, Steven Chamberland, and Samuel Pierre from the École Polytechnique de Montréal, Montréal Québec, Canada. A local search heuristic for the global planning of UMTS networks. In Seizo Onoe, Mohsen Guizani, Hsiao-Hwa Chen, and Mamoru Sawahashi, editors, *Proceedings of the International Conference on Wireless Communications and Mobile Computing, IWCMC 2006*, pages 1405–1410, July 3–6 2006, Vancouver, British Columbia, Canada. ACM, New York, NY, USA. ISBN 1-59593-306-9. doi:10.1145/1143549.1143831. SESSION: R2-D: general symposium. Online available at <http://doi.acm.org/10.1145/1143549.1143831> [accessed 2008-08-01].
 - [1454] Barry D. Hughes. *Random Walks and Random Environments: Volume 1: Random Walks*. Oxford University Press, USA, May 16, 1995. ISBN 0198537883, 978-0198537885.
 - [1455] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, 3 edition, 1968. ISBN 0471257087, 978-0471257080.
 - [1456] Sticky Fingaz. Throw Ya Gunz. In Onyx, editor, *Bacdafucup*, Chyskillz, producer. Def Jam, March 30, 1993. Catalog Number: O 53302.
 - [1457] John H. Gillespie. Molecular Evolution Over the Mutational Landscape. *Evolution*, 38(5):1116–1129, September 1984. doi:10.2307/2408444.
 - [1458] John Maynard Smith from the School of Biological Sciences, University of Sussex. Natural Selection and the Concept of a Protein Space. *Nature*, 225(5232):563–564, February 7, 1970. ISSN 0028-0836. doi:10.1038/225563a0. Received November 7, 1969.
 - [1459] Herb Sutter and James Larus. Software and the concurrency revolution. *Queue*, 3(7):54–62, 2005. ISSN 1542-7730. Online available at <http://research.microsoft.com/~larus/Papers/queue01.pdf> and <http://portal.acm.org/citation.cfm?id=1095421> [accessed 2007-08-13].
 - [1460] Omer Berkman, Dany Breslauer, Zvi Galil, Baruch Schieber, and Uzi Vishkin. Highly parallelizable problems. In *STOC’89: Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 309–319, 1989, Seattle, Washington, United States. ACM Press, New York, NY, USA. ISBN 0-89791-307-8. doi:10.1145/73007.73036. Online available at <http://doi.acm.org/10.1145/73007.73036> and <http://libra.msra.cn/paperDetail.aspx?id=907688> [accessed 2007-08-13].
 - [1461] Valmir C. Barbosa. *An Introduction to Distributed Algorithms*. The MIT Press, September 1996. ISBN 978-0262024129.
 - [1462] Gerard Tel. *Introduction to Distributed Algorithms*. Cambridge University Press, second edition, January 2004. ISBN 978-0521794831.

- [1463] Enrique Alba and Marco Tomassini. Parallelism and evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 6(5):443–462, October 2002. doi:10.1109/TEVC.2002.800880.
- [1464] Erick Cant’u-Paz. *Designing efficient and accurate parallel genetic algorithms*. PhD thesis, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, July 1999. Also IlliGAL report 99017. Online available at <http://citeseer.ist.psu.edu/cantu-paz99designing.html> [accessed 2008-04-06].
- [1465] Kay Chen Tan, Arthur Tay, and Ji Cai. Design and implementation of a distributed evolutionary computing software. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 33(3):325–338, August 2003. doi:10.1109/TSMCC.2003.817359.
- [1466] Reiko Tanese. Distributed Genetic Algorithms. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 434–439, 1989. In proceedings [652].
- [1467] Heinz Mühlenbein. Evolution in Time and Space – The Parallel Genetic Algorithm. In *Foundations of Genetic Algorithms 1*, pages 316–337, 1990. In proceedings [188]. Online available at <http://muehlenbein.org/parallel.PDF> and <http://citeseer.ist.psu.edu/11201.html> [accessed 2007-11-30].
- [1468] Alessandro Bollini and Marco Piastra. Distributed and Persistent Evolutionary Algorithms: a Design Pattern. In *Genetic Programming, Proceedings of EuroGP’99*, pages 173–183, 1999. In proceedings [910].
- [1469] Dean M. Tullsen, Susan Eggers, and Henry M. Levy. Simultaneous Multithreading: Maximizing On-Chip Parallelism. In *Proceedings of the 22th Annual International Symposium on Computer Architecture*, pages 392–403. ACM Press, June 1995, Santa Margherita Ligure, Italy. Online available at <http://www.cs.washington.edu/research/smt/papers/ISCA95.ps> and <http://citeseer.ist.psu.edu/tullsen95simultaneous.html> [accessed 2007-09-11].
- [1470] Kunle Olukotun, Basem A. Nayfeh, Lance Hammond, Ken Wilson, and Kunyung Chang. The case for a single-chip multiprocessor. In *ASPLOS-VII: Proceedings of the seventh international conference on Architectural support for programming languages and operating systems*, pages 2–11, 1996, Cambridge, Massachusetts, United States. ACM Press, New York, NY, USA. ISBN 0-89791-767-7. Online available at http://ogun.stanford.edu/~kunle/publications/hydra_ASPLoS_VII.pdf and <http://doi.acm.org/10.1145/237090.237140> [accessed 2007-09-11].
- [1471] Geoff Koch from the Science Writer, Michigan State University. Discovering Multi-Core: Extending the Benefits of Moore’s Law. *Technology@Intel Magazine*, July 2005. Online available at <http://www.intel.com/technology/magazine/computing/multi-core-0705.pdf> [accessed 2007-09-11]. See also <http://www.intel.com/technology/magazine/> [accessed 2007-09-11].
- [1472] Michael Singer. Intel First to Ship Dual Core x86 chip. *internetnews.com*, April 2005. Online available at <http://www.internetnews.com/ent-news/article.php/3496926> [accessed 2007-07-16]. See also <http://www.internetnews.com/> [accessed 2007-09-11].
- [1473] Erick Cant’u-Paz. A Summary of Research on Parallel Genetic Algorithms. IlliGAL report 95007, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, 1995. Online available at <http://citeseer.ist.psu.edu/27505.html> and <ftp://ftp-illigal.ge.uiuc.edu/pub/papers/IlliGALs/95007.ps.Z> [accessed 2007-08-13].
- [1474] David A. Van Veldhuizen, Jesse B. Zydallis, and Gary B. Lamont. Issues in parallelizing multiobjective evolutionary algorithms for real world applications. In *SAC’02: Proceedings of the 2002 ACM symposium on Applied computing*, pages 595–602, 2002, Madrid, Spain. ACM Press, New York, NY, USA. ISBN 1-58113-445-2. doi:<http://doi.acm.org/10.1145/508791.508906>. Online available at <http://doi.acm.org/10.1145/508791.508906> [accessed 2007-08-14].

- [1475] Kai Xu, Sushil J. Louis, and Roberto C. Mancini. A scalable parallel genetic algorithm for x-ray spectroscopic analysis. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 811–816, 2005. doi:<http://doi.acm.org/10.1145/1068009.1068145>. In proceedings [554]. Online available at <http://doi.acm.org/10.1145/1068009.1068145> [accessed 2007-08-14].
- [1476] Marc Dubreuil, Christian Gagné, and Marc Parizeau. Analysis of a master-slave architecture for distributed evolutionary computations. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 36(1):229–235, February 2006. doi:10.1109/TSMCB.2005.856724. Online available at <http://vision.gel.ulaval.ca/~parizeau/publications/smc06.pdf> [accessed 2008-04-06].
- [1477] Sébastien Cahon, Nordine Melab, and El-Ghazali Talbi. ParadisEO: A Framework for the Reusable Design of Parallel and Distributed Metaheuristics. *Journal of Heuristics*, 10(3):357–380, May 2004. ISSN 1381-1231 (Print) 1572-9397 (Online). doi:10.1023/B:HEUR.0000026900.92269.ec.
- [1478] Zbigniew Skolicki and Kenneth Alan De Jong. The influence of migration sizes and intervals on island models. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1295–1302, 2005. doi:<http://doi.acm.org/10.1145/1068009.1068219>. In proceedings [554]. Online available at <http://portal.acm.org/citation.cfm?id=1068009.1068219> [accessed 2007-08-14].
- [1479] Zbigniew Skolicki. An analysis of island models in evolutionary computation. In *Workshop Proceedings of Genetic and Evolutionary Computation Conference, GECCO 2005*, pages 386–389, 2005. In proceedings [555]. Online available at <http://cs.gmu.edu/~eclab/papers/skolicki05analysis.pdf> and <http://portal.acm.org/citation.cfm?id=1102256.1102343> [accessed 2007-08-14].
- [1480] Martina Gorges-Schleuter. Explicit Parallelism of Genetic Algorithms through Population Structures. In *PPSN I: Proceedings of the 1st Workshop on Parallel Problem Solving from Nature*, pages 150–159, 1991. ISBN 3-540-54148-9. In proceedings [592].
- [1481] Naoyoshi Toshine, Nobuyuki Iwauchi, Shin'ichi Wakabayashi, T. Koide, and I. Nishimura from the Faculty of Engineering, Hiroshima University. A parallel genetic algorithm with adaptive adjustment of genetic parameters. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pages 679–686, 2001. In proceedings [568].
- [1482] L. Darrell Whitley and Timothy Starkweather. GENITOR II: A Distributed Genetic Algorithm. *Journal of Experimental & Theoretical Artificial Intelligence*, 2(3):189–214, July 1990. ISSN 0952-813X. doi:10.1080/09528139008953723.
- [1483] B. Paechter, Thomas Bäck, Marc Schoenauer, M. Sebag, Agoston E. Eiben, J. J. Merelo, and T. C. Fogarty. A Distributed Resource Evolutionary Algorithm Machine (DREAM). In *Proceedings of the Congress on Evolutionary Computation*, volume 2, pages 951–958, 2000. doi:10.1109/CEC.2000.870746. In proceedings [501].
- [1484] Maribel G. Arenas, Pierre Collet, Agoston E. Eiben, Márk Jelasity, Juan J. Merelo, Ben Paechter, Mike Preuß, and Marc Schoenauer. A Framework for Distributed Evolutionary Algorithms. In *Proceedings of the 7th International Conference on Parallel Problem Solving from Nature – PPSN VII*, pages 665–675, 2002. In proceedings [586]. Online available at <http://citeseer.ist.psu.edu/arenas02framework.html> [accessed 2008-04-06].
- [1485] Fuey Sian Chong and William B. Langdon. Java based Distributed Genetic Programming on the Internet. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, page 1229, 1999. In proceedings [572]. Online available at <http://libra.msra.cn/paperDetail.aspx?id=589775> and <http://citeseer.ist.psu.edu/chong99java.html> [accessed 2007-08-14].
- [1486] Hisao Ishibuchi, Tadashi Yoshida, and Tadahiko Murata. Balance Between Genetic Search And Local Search In Hybrid Evolutionary Multi-criterion Optimization Al-

- gorithms. In *GECCO'02: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1301–1308, 2002. In proceedings [563]. Online available at <http://citeseer.ist.psu.edu/632937.html> [accessed 2007-09-10].
- [1487] Miguel Alabau, Lhassane Idoumghar, and René Schott. New hybrid genetic algorithms for the frequency assignment problem. In *Proceedings of the 13th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'01)*, page 136 ff. IEEE Computer Society, November 7–9, 2001, Dallas, TX, USA. ISBN 0-7695-1417-0.
- [1488] Vincent Bachelet and El-Ghazali Talbi. A Parallel Co-evolutionary Metaheuristic. In *IPDPS'00: Proceedings of the 15 IPDPS 2000 Workshops on Parallel and Distributed Processing*, pages 628–635, 2000. Springer-Verlag, London, UK. ISBN 3-540-67442-X. Online available at <http://ipdps.cc.gatech.edu/2000/biosp3/18000629.pdf> and <http://citeseer.ist.psu.edu/bachelet00parallel.html> [accessed 2007-09-10].
- [1489] Xin Yao from the Computer Sciences Laboratory, Research School of Physical Sciences, The Australian National University. Optimization by genetic annealing. In M. Jabri, editor, *Proceedings of Second Australian Conference on Neural Networks*, pages 94–97, 1991, Sydney, Australia. Online available at <http://citeseer.ist.psu.edu/yao91optimization.html> and <ftp://www.cs.adfa.edu.au/pub/xin/acnn91.ps.Z> [accessed 2007-09-10].
- [1490] Derek M. Johnson, Ankur M. Teredesai, and Robert T. Saltarelli from the Computer Science Department, Rochester Institute of Technology, Rochester NY 14623, USA. Genetic Programming in Wireless Sensor Networks. In *Genetic Programming*, pages 96–107, 2005. doi:10.1007/b107383. In proceedings [904]. Online available at <http://www.cs.rit.edu/~amt/pubs/EuroGP05FinalTeredesai.pdf> [accessed 2008-06-17].
- [1491] L. Darrell Whitley. Cellular Genetic Algorithms. In *Proceedings of the 5th International Conference on Genetic Algorithms*, page 658, 1993. In proceedings [748].
- [1492] Bernard Manderick and Piet Spiessens. Fine-grained parallel genetic algorithms. In *Proceedings of the third international conference on Genetic algorithms*, pages 428–433, 1989. In proceedings [652].
- [1493] W. Daniel Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D: Nonlinear Phenomena*, 42(1-3):228–234, June 1990. doi:10.1016/0167-2789(90)90076-2. Provided by the SAO/NASA Astrophysics Data System.
- [1494] Yuval Davidor. A Naturally Occurring Niche-Species Phenomenon – The Model and First Results. In *Proceedings of the 4th International Conference on Genetic Algorithms*, pages 257–263, 1989. In proceedings [749].
- [1495] Robert J. Collins and David R. Jefferson. Selection in Massively Parallel Genetic Algorithms. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 249–256, 1991. In proceedings [749]. Online available at <http://citeseer.ist.psu.edu/349839.html> [accessed 2007-11-30].
- [1496] Joel N. Morse. Reducing the size of the nondominated set: Pruning by clustering. *Computers & Operations Research*, 7(1–2):55–66, 1980. doi:10.1016/0305-0548(80)90014-3.
- [1497] Heidi A. Taboada, Fatema Baheranwala, David W. Coit, and Naruemon Wattanapongsakorn. Practical Solutions for Multi-Objective Optimization: An Application to System Reliability Design Problems. *Reliability Engineering & System Safety*, 92(3):314–322, March 2007. doi:10.1016/j.res.2006.04.014. Online available at http://www.rci.rutgers.edu/~coit/RESS_2007.pdf [accessed 2007-09-10].
- [1498] David W. Coit and Fatema Baheranwala. Solution of stochastic multi-objective system reliability design problems using genetic algorithms. In *Proceedings of the European safety and reliability conference (ESREL)*, pages 391–398, June 2005, Gdansk, Poland. Online available at http://www.engr.rutgers.edu/ie/research/working_paper.html [accessed 2007-07-29].

- [1499] Heidi A. Taboada and David W. Coit. Data clustering of solutions for multiple objective system reliability optimization problems. *Quality Technology and Quantitative Management Journal*, 4(2):35–54, June 2007. Online available at http://www.soe.rutgers.edu/ie/research/working_paper/paper05-019.pdf [accessed 2007-09-10].
- [1500] Jonathan E. Fieldsend, Richard M. Everson, and Sameer Singh. Using Unconstrained Elite Archives for Multi-Objective Optimisation. *IEEE Transactions on Evolutionary Computing*, 7(3):305–323, 2003. This document supersedes the manuscript [1949]. Online available at <http://citeseer.ist.psu.edu/543052.html> and <http://www.lania.mx/~ccoello/EM00/fieldsend03.pdf.gz> [accessed 2007-08-25].
- [1501] Joshua D. Knowles and David W. Corne. Properties of an adaptive archiving algorithm for storing nondominated vectors. *IEEE Transactions on Evolutionary Computation*, 7:100–116, April 2003.
- [1502] George Edward Pelham Box, J. Stuart Hunter, and William G. Hunter. *Statistics for Experimenters: Design, Innovation, and Discovery*. Wiley & Son, 2nd edition, May 2005. ISBN 0-471-71813-0, 978-0471718130.
- [1503] Frank Yates. *The Design and Analysis of Factorial Experiments*. Imperial Bureau of Soil Science, Commonwealth Agricultural Bureaux, Harpenden, England, UK, 1937. ISBN 0851982204, 978-0851982205. ASIN: B00086SIZ0. Technical Communication No. 35.
- [1504] Sir Ronald Aylmer Fisher. The Arrangement of Field Experiments. *Journal of the Ministry of Agriculture of Great Britain*, 33:503–513, 1926. Online available at <http://digital.library.adelaide.edu.au/coll/special//fisher/48.pdf> [accessed 2008-08-07].
- [1505] Sean Luke and Liviu Panait. A survey and comparison of tree generation algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 81–88, 2001. In proceedings [568]. Online available at <http://citeseer.ist.psu.edu/luke01survey.html> and <http://www.cs.gmu.edu/~sean/papers/treegenalgs.pdf> [accessed 2007-07-28].
- [1506] Pu Liu, Francis Lau, Michael J. Lewis, and Cho li Wang. A New Asynchronous Parallel Evolutionary Algorithm for Function Optimization. In *Proceedings of the 7th International Conference on Parallel Problem Solving from Nature*, pages 401–410, 2002. In proceedings [586]. Online available at [accessed 2008-04-06].
- [1507] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evolutionary Computation*, 8(2): 173–195, 2000. Online available at <http://sci2s.ugr.es/docencia/cursosMieres/EC-2000-Comparison.pdf> and <http://citeseer.comp.nus.edu.sg/362080.html> [accessed 2008-04-06].
- [1508] Robin C. Purshouse and Peter J. Fleming. The Multi-Objective Genetic Algorithm Applied to Benchmark Problems – An Analysis. Research Report 796, Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, S1 3JD, UK, August 2001. Online available at <http://citeseer.ist.psu.edu/499210.html> and <http://www.lania.mx/~ccoello/EM00/purshouse01.pdf.gz> [accessed 2008-04-16].
- [1509] Thomas Jansen and Ingo Wegener from the FB 4, LS 2, Universitt Dortmund, 44221 Dortmund, Germany. A comparison of simulated annealing with a simple evolutionary algorithm on pseudo-boolean functions of unitation. *Theoretical Computer Science*, 386(1-2):73–93, October 2007. ISSN 0304-3975. doi:10.1016/j.tcs.2007.06.003. Online available at <http://dx.doi.org/10.1016/j.tcs.2007.06.003> [accessed 2008-04-07].
- [1510] Edmund Burke, Steven Matt Gustafson, Graham Kendall, and Natalio Krasnogor. Advanced Population Diversity Measures in Genetic Programming. In *Proceedings of the 7th International Conference on Parallel Problem Solving from Nature - PPSN*

- VII, page 341 ff., 2002. In proceedings [586]. Online available at <http://citeseer.ist.psu.edu/529057.html> [accessed 2007-09-07].
- [1511] P. N. Suganthan, N. Hansen, J. J. Liang, Kalyanmoy Deb, Y. P. Chen, Anne Auger, and S. Tiwari. Problem Definitions and Evaluation Criteria for the CEC 2005 Special Session on Real-Parameter Optimization. KanGAL Report 2005005, Kanpur Genetic Algorithms Laboratory, KanGAL, Indian Institute of Technology Kanpur, India, May 2005. Online available at <http://www.iitk.ac.in/kangal/papers/k2005005.pdf> [accessed 2007-10-07]. See also [1950, 496].
- [1512] Stuart Alan Kauffman and Simon Asher Levin from the Department of Biochemistry and Biophysics, University of Pennsylvania School of Medicine, Philadelphia 19104, USA. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45, September 7, 1987. ISSN 0022-5193.
- [1513] Stuart Alan Kauffman and Edward D. Weinberger. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology*, 141(2):211–245, November 21, 1989. Online available at [http://dx.doi.org/10.1016/S0022-5193\(89\)80019-0](http://dx.doi.org/10.1016/S0022-5193(89)80019-0) [accessed 2007-10-14].
- [1514] Edward D. Weinberger. Local properties of Kauffman’s NK model, a tuneably rugged energy landscape. *Physical Review A*, 44(10):6399–6413, 1991.
- [1515] Walter Fontana, Peter F. Stadler, Erich G. Bornberg-Bauer, Thomas Griesmacher, Ivo L. Hofacker, Manfred Tacker, Pedro Tarazona, Edward D. Weinberger, and Peter Schuster. RNA folding and combinatorial landscapes. *Physical Review E*, 47(3):2083–2099, March 1993. doi:10.1103/PhysRevE.47.2083. Online available at <http://dx.doi.org/10.1103/PhysRevE.47.2083> and <http://citeseer.ist.psu.edu/94823.html> [accessed 2007-11-27].
- [1516] Michael Defoin Platel, Sébastien Vérel, Manuel Clergue, and Philippe Collard. From Royal Road to Epistatic Road for Variable Length Evolution Algorithm. In *Evolution Artificielle, 6th International Conference*, pages 3–14, 2003. doi:10.1007/b96080. In proceedings [512]. Online available at <http://arxiv.org/abs/0707.0548> and <http://www.i3s.unice.fr/~verel/publi/ea03-fromRRtoER.pdf> [accessed 2007-12-01].
- [1517] Henrik Flyvbjerg and Benny Lautrup from the CONNECT, The Nils Bohr Institute, Blegdamsvej 17, DK-2100, Copenhagen, Denmark. Evolution in a Rugged Fitness Landscape. *Physical Review A*, 46(10):6714–6723, November 15, 1992. doi:10.1103/PhysRevA.46.6714. Online available at <http://citeseer.ist.psu.edu/26647.html> and http://prola.aps.org/abstract/PRA/v46/i10/p6714_1 [accessed 2007-08-14].
- [1518] Edward D. Weinberger. $\mathcal{NP}$ Completeness of Kauffman’s N-k Model, A Tuneable Rugged Fitness Landscape. Working Papers 96-02-003, Santa Fe Institute, February 1996. Online available at <http://www.santafe.edu/research/publications/workingpapers/96-02-003.ps> [accessed 2007-08-24].
- [1519] Richard K. Thompson and Alden H. Wright. Additively Decomposable Fitness Functions. Technical Report, Computer Science Department, The University of Montana, Missoula, MT 59812-1008, USA, 1996. Online available at <http://citeseer.ist.psu.edu/thompson96additively.html> and http://www.cs.umt.edu/u/wright/papers/add_decomp.ps.gz [accessed 2007-11-27].
- [1520] Nicholas Geard. An exploration of NK landscapes with neutrality. Technical Report 14213, School of Information Technology and Electrical Engineering, University of Queensland, St Lucia Q 4072, Australia, October 2001. Supervisor: Janet Wiles. Online available at http://eprints.ecs.soton.ac.uk/14213/1/ng_thesis.pdf [accessed 2008-07-01].
- [1521] Nicholas Geard, Janet Wiles, Jennifer Hallinan, Bradley Tonkes, and Ben Skellett from the The University of Queensland, Brisbane Q 4072, Australia. A Comparison of Neutral Landscapes – NK, NKp and NKq. In *Congress on Evolutionary Computation (CEC2002)*, pages 205–210, 2002. In proceedings [499]. Online available

- at <http://citeseer.ist.psu.edu/506371.html> and <http://eprints.ecs.soton.ac.uk/14208/1/cec1-comp.pdf> [accessed 2008-06-01].
- [1522] Christian M. Reidys and Peter F. Stadler. Neutrality in fitness landscapes. *Applied Mathematics and Computation*, 117(2-3):321-350, January 25, 2001. ISSN 0096-3003. doi:10.1016/S0096-3003(99)00166-6. Online available at [http://dx.doi.org/10.1016/S0096-3003\(99\)00166-6](http://dx.doi.org/10.1016/S0096-3003(99)00166-6) and <http://citeseer.ist.psu.edu/reidys98neutrality.html> [accessed 2008-06-01].
 - [1523] M. E. J. Newman and R. Engelhardt. Effect of neutral selection on the evolution of molecular species. *Proceedings of the Royal Society of London B (Biological Sciences)*, 256(1403):1333-1338, July 22, 1998. ISSN 0962-8452 (Print) 1471-2954 (Online). doi:10.1098/rspb.1998.0438. Online available at <http://citeseer.ist.psu.edu/202669.html> and <http://journals.royalsociety.org/content/5ttwagyubu88bang/fulltext.pdf> [accessed 2008-06-01].
 - [1524] José Lobo, John H. Miller, and Walter Fontana. Neutrality in technological landscapes. Santa fe working paper, Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA, January 26, 2004. Online available at <http://fontana.med.harvard.edu/www/Documents/WF/Papers/tech.neut.pdf> and <http://citeseer.ist.psu.edu/lobo04neutrality.html> [accessed 2008-06-06].
 - [1525] Philip Warren Anderson from the Bell Laboratories, Murray Hill, New Jersey 07974; and Joseph Henry Laboratories, Princeton University, Princeton, New Jersey 08544, USA. Suggested Model for Prebiotic Evolution: The Use of Chaos. *Proceedings of the National Academy of Science of the United States of America (PNAS)*, 80(11):3386-3390, June 1, 1983. Online available at <http://www.pnas.org/cgi/reprint/80/11/3386.pdf> and <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=394048> [accessed 2008-07-02].
 - [1526] Daniel S. Rokhsar, Philip Warren Anderson, and Daniel L. Stein. Self-organization in prebiological systems: Simulations of a model for the origin of genetic information. *Journal of Molecular Evolution*, 23(2):119-126, June 1986. ISSN 0022-2844 (print), 1432-1432 (online). doi:10.1007/BF02099906. Online available at <http://www.springerlink.com/content/m268540162062735/fulltext.pdf> [accessed 2008-07-02].
 - [1527] Constantino Tsallis and Ricardo Ferreira from the Centro Brasileiro de Pesquisas Físicas/CNPq, Rua Dr. Xavier Sigaud 150, 22290 Rio de Janeiro, RJ, Brazil. On the origin of self-replicating information-containing polymers from oligomeric mixtures. *Physics Letters A*, 99(9):461-463, December 26, 1983. doi:10.1016/0375-9601(83)90958-1. Online available at [http://dx.doi.org/10.1016/0375-9601\(83\)90958-1](http://dx.doi.org/10.1016/0375-9601(83)90958-1) [accessed 2008-07-01].
 - [1528] Ciro Amitrano, Luca Peliti, and M. Saber from the Dipartimento di Scienze Fisiche, Università di Napoli, Italy. Population dynamics in a spin-glass model of chemical evolution. *Journal of Molecular Evolution*, 29(6):513-525, December 1989. ISSN 0022-2844 (Print) 1432-1432 (Online). doi:10.1007/BF02602923. Online available at <http://www.springerlink.com/content/n172u715884w0632/fulltext.pdf> [accessed 2008-07-02].
 - [1529] Edward D. Weinberger and Peter F. Stadler. Why some fitness landscapes are fractal. *Journal of Theoretical Biology*, 163(2):255-275., July 21, 1993. doi:10.1006/jtbi.1993.1120. Online available at <http://www.tbi.univie.ac.at/papers/Abstracts/93-01-01.ps.gz> and <http://dx.doi.org/10.1006/jtbi.1993.1120> [accessed 2008-06-01].
 - [1530] Kalyanmoy Deb and David E. Goldberg. Analyzing deception in trap functions. In *Foundations of Genetic Algorithms 2*, pages 93-108, 1993. In proceedings [738].
 - [1531] David H. Ackley. *A connectionist machine for genetic hillclimbing*, volume 28 of *The Springer International Series in Engineering and Computer Science*. Kluwer

- Academic Publishers, Norwell, MA, USA, 1987. ISBN 0-89838-236-X, 978-0-89838-236-5.
- [1532] Terry Jones. A Description of Holland's Royal Road Function. *Evolutionary Computation*, 2(4):411–417, 1994. See also [1951].
- [1533] Stephanie Forrest and Melanie Mitchell. Relative Building-Block Fitness and the Building-Block Hypothesis. In *Foundations of Genetic Algorithms 2*, pages 109–126, 1992. In proceedings [738]. Online available at <http://web.cecs.pdx.edu/~mm/Forrest-Mitchell-FOGA.pdf> and <http://citeseer.ist.psu.edu/39682.html> [accessed 2007-08-13].
- [1534] R. J. Quick, Victor J. Rayward-Smith, and George D. Smith from the University of East Anglia, Norwich, UK. The Royal Road functions: description, intent and experimentation. In *Evolutionary Computing – Selected Papers from the AISB Workshop*, volume 1143/1996 of *Lecture Notes in Computer Science (LNCS)*, pages 223–235. Springer Berlin/Heidelberg, April 12, 1996, Brighton, U.K. ISBN 978-3-540-61749-5. doi:10.1007/BFb0032786. Partly online available at <http://books.google.de/books?id=cJi2MRos9cMC> [accessed 2008-04-09].
- [1535] Michael Defoin Platel, Manuel Clergue, and Philippe Collard. Maximum Homologous Crossover for Linear Genetic Programming. In *Genetic Programming: 6th European Conference*, pages 29–48, 2003. In proceedings [906]. Online available at http://www.i3s.unice.fr/~clergue/siteCV/publi/eurogp_03.pdf [accessed 2007-12-01].
- [1536] William F. Punch, Douglas Zongker, and Erik D. Goodman. The Royal Tree Problem, a Benchmark for Single and Multiple Population Genetic Programming. In *Advances in Genetic Programming 2*, pages 299–316. MIT Press, 1996. In collection [929]. Online available at <http://citeseer.ist.psu.edu/147908.html> [accessed 2007-10-14].
- [1537] Tobias Storch and Ingo Wegener from the Department of Computer Science, University of Dortmund, 44221 Dortmund, Germany. Real Royal Road Functions for Constant Population Size. In *Genetic and Evolutionary Computation – GECCO 2003*, pages 1406–1417, 2003. doi:10.1007/3-540-45110-2_14. In proceedings [562]. See also [1538, 1539].
- [1538] Tobias Storch and Ingo Wegener from the University of Dortmund, Department of Computer Science/XI, 44221 Dortmund, Germany. Real Royal Road Functions for Constant Population Size. Computational Intelligence CI-167/04, Design and Management of Complex Technical Processes and Systems by means of Computational Intelligence Methods, Collaborative Research Center 531, University of Dortmund, February 2004. Secretary of the SFB 531. Online available at <http://hdl.handle.net/2003/5456> [accessed 2008-07-22]. See also [1537, 1539].
- [1539] Tobias Storch and Ingo Wegener from the Department of Computer Science, University of Dortmund, 44221 Dortmund, Germany. Real Royal Road Functions for Constant Population Size. *Theoretical Computer Science*, 320(1):123–134, June 12, 2004. ISSN 0304-3975. doi:10.1016/j.tcs.2004.03.047. Online available at <http://dx.doi.org/10.1016/j.tcs.2004.03.047> [accessed 2008-07-22]. See also [1537, 1538].
- [1540] Heinz Mühlenbein and Dirk Schlierkamp-Voosen from the Fraunhofer GMD P.O. 1316, D-5205 Sankt Augustin 1, Germany. Predictive models for the breeder genetic algorithm I: Continuous parameter optimization. *Evolutionary Computations*, 1(1): 25–49, Spring 1993. ISSN 1063-6560, 2049097. Online available at <http://citeseer.ist.psu.edu/mtihlenbein93predictive.html> [accessed 2008-07-21].
- [1541] Dirk Thierens and David E. Goldberg. Convergence Models of Genetic Algorithm Selection Schemes. In *PPSN III: Proceedings of the International Conference on Evolutionary Computation. The Third Conference on Parallel Problem Solving from Nature*, pages 119–129, 1994. In proceedings [590]. Online available at <http://www.cs.uu.nl/groups/DSS/publications/convergence/convMdl.ps> [accessed 2008-07-21].
- [1542] Thomas Bäck. Generalized Convergence Models for Tournament- and (μ, λ) -Selection. In *Proceedings of the 6th International Conference on Genetic Algorithms*,

- pages 2–8, 1995. In proceedings [747]. Online available at <http://citeseer.ist.psu.edu/1463.html> [accessed 2008-07-21].
- [1543] Dominic Wilson and Devinder Kaur from the Electrical Engineering and Computer Science, University of Toledo, Toledo, Ohio 43606, USA. Using Quotient Graphs to Model Neutrality in Evolutionary Search. In *Genetic and Evolutionary Computation Conference*, pages 2233–2238, 2008. In proceedings [550].
- [1544] Kumara Sastry and David E. Goldberg. Let’s get ready to rumble redux: crossover versus mutation head to head on exponentially scaled problems. In *GECCO’07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 1380–1387, 2007. doi:10.1145/1276958.1277215. In proceedings [551]. Online available at <http://doi.acm.org/10.1145/1276958.1277215> [accessed 2008-07-21]. See also [1545].
- [1545] Kumara Sastry and David E. Goldberg. Let’s get ready to rumble redux: Crossover versus mutation head to head on exponentially scaled problems. IlliGAL Report 2007005, Illinois Genetic Algorithms Laboratory (IlliGAL), University of Illinois, Urbana-Champaign, Urbana IL, USA, February 11, 2007. Online available at <http://www.illigal.uiuc.edu/pub/papers/IlliGALs/2007005.pdf> [accessed 2008-07-21]. See also [1544].
- [1546] Thomas Weise, Stefan Niemczyk, Hendrik Skubch, Roland Reichle, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. A Tunable Model for Multi-Objective, Epistatic, Rugged, and Neutral Fitness Landscapes. In *Proceedings of Genetic and Evolutionary Computation Conference, GECCO 2008*, pages 795–802, 2008. In proceedings [550]. Online available at <http://www.it-weise.de/documents/index.html#WNSRG2008GECCO> [accessed 2008-08-12].
- [1547] Stefan Niemczyk. Ein Modellproblem mit einstellbarer Schwierigkeit zur Evaluierung Evolutionärer Algorithmen. Master’s thesis, Universität Kassel, Fachgebiet Elektrotechnik/Informatik, Fachbereich Verteilte Systeme, Wilhelmshöher Allee 73, 34121 Kassel, Deutschland, May 5, 2008. Diplom I Arbeit (equiv. Bachelor). Advisor: Thomas Weise. Supervisor: Kurt Geihs. Committee: Albert Zündorf. Online available at <http://www.it-weise.de/documents/index.html#N2008MPMES> [accessed 2008-08-12].
- [1548] Robert J. Collins and David R. Jefferson. Representations for artificial organisms. In *From animals to animats: Proceedings of the first international conference on simulation of adaptive behavior*, pages 382–390, 1990. In proceedings [1925].
- [1549] David R. Jefferson, Robert J. Collins, Claus Cooper, Michael Dyer, Margot Flowers, Richard Korf, Charles Tayler, and Alan Wang. Evolution as a theme in artificial life: The Genesys/Tracker system. In *Artificial Life II*, pages 549–578, 1990. In proceedings [1924].
- [1550] Robert J. Collins and David R. Jefferson. AntFarm: Towards Simulated Evolution. In *Artificial Life II*, pages 579–601, 1992. In proceedings [1924]. Online available at <http://citeseer.ist.psu.edu/collins91antfarm.html> and <http://en.scientificcommons.org/108082> [accessed 2007-08-05].
- [1551] Kim Harries and Peter W. H. Smith from the Department of Computer Science, City University, Northampton Square, London EC1V 0HB, United Kingdom. Exploring Alternative Operators and Search Strategies in Genetic Programming. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 147–155, 1997. In proceedings [915]. Online available at <http://citeseer.ist.psu.edu/harries97exploring.html> and http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/harries.gp97_paper.ps.gz [accessed 2008-04-12].
- [1552] Sean Luke and Lee Spector. A Comparison of Crossover and Mutation in Genetic Programming. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 240–248, 1997. In proceedings [915]. Online avail-

- able at <http://citeseer.ist.psu.edu/146629.html> and <http://hampshire.edu/lspector/pubs/comparison.pdf> [accessed 2008-04-12].
- [1553] Ibrahim Kuscü. Evolving a Generalised Behavior: Artificial Ant Problem Revisited. In *7th Annual Conference on Evolutionary Programming*, page 799 ff., 1998. ISBN 3-540-64891-7. In proceedings [1186].
 - [1554] Kumar Chellapilla. Evolutionary Programming with Tree Mutations: Evolving Computer Programs without Crossover. In *Genetic Programming 1997: Proceedings of the Second Annual Conference GP-97*, pages 432–438, 1997. In proceedings [915].
 - [1555] Takuya Ito, Hitoshi Iba, and Satoshi Sato. Non-destructive depth-dependent crossover for genetic programming. In *Genetic Programming*, pages 71–82, 1998. doi:10.1007/BFb0055929. In proceedings [911].
 - [1556] William B. Langdon and Riccardo Poli from the School of Computer Science, University of Birmingham, Birmingham B15 2TT, UK. Better Trained Ants for Genetic Programming. Technical Report CSRP-98-12, University of Birmingham, School of Computer Science, April 2, 1998. Online available at <http://citeseer.ist.psu.edu/105696.html> and <ftp://ftp.cs.bham.ac.uk/pub/tech-reports/1998/CSRP-98-12.ps.gz> [accessed 2007-08-05].
 - [1557] Clemens Frey. *Virtual Ecosystems – Evolutionary and Genetic Programming from the perspective of modern means of ecosystem-modelling*, volume 93 of *Bayreuth Forum Ecology*. Institute for Terrestrial Ecosystems, Bayreuth, Bayreuth, Germany, 2002.
 - [1558] Christopher G. Langdon, editor. *Artificial Life: The Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume 6 of *Santa Fe Institute Studies in the Science of Complexity*, September 1987, Los Alamos, New Mexico, USA. Addison-Wesley Publishing Company, Reading, MA, USA / Redwood City, CA, USA. ISBN 0201093561, 978-0201093568. Published in January 1989.
 - [1559] John R. Koza. Genetically breeding populations of computer programs to solve problems in artificial intelligence. In *Proceedings of the Second International Conference on Tools for AI, Herndon, Virginia, USA*, pages 819–827. IEEE Computer Society Press, Los Alamitos, CA, USA, 6-9 1990. Online available at <http://citeseer.ist.psu.edu/koza90genetically.html> and <http://www.lania.mx/~ccoello/koza90.ps.gz> [accessed 2007-09-09].
 - [1560] Brian Connolly. Genetic Algorithms – Survival of the Fittest: Natural Selection with Windows Forms. *MSDN Magazin*, August 2004. Online available at [http://msdn.microsoft.com/de-de/magazine/cc163934\(en-us\).aspx](http://msdn.microsoft.com/de-de/magazine/cc163934(en-us).aspx) and 2008-06-24 [accessed 2008-06-24].
 - [1561] Euclid of Alexandria, editor. *Stoicheia (Elements)*. Euclid of Alexandria, Alexandria, 300 BC. A series consisting of 13 books, to which two books were added lateron. Online available at <http://aleph0.clarku.edu/~djoyce/java/elements/elements.html> and <http://www.gutenberg.org/etext/21076> [accessed 2007-10-05] (in english). Online available at <http://farside.ph.utexas.edu/euclid.html> and <http://en.wikipedia.org/wiki/Image:Euclid-Elements.pdf> [accessed 2007-10-05] (bilingual/-greek). See also [1562] and http://en.wikipedia.org/wiki/Euclid%27s_Elements [accessed 2007-10-05].
 - [1562] Thomas L. Heath. *The Thirteen Books of Euclid's Elements*. Dover Publications, New York, second edition, 1956. ISBN 0-486-60088-2 (vol. 1), 0-486-60089-0 (vol. 2), 0-486-60090-4 (vol. 3). Three volumes. Original publication: Cambridge University Press, 1925. Euclid's original: [1561].
 - [1563] Philipp Andreas Baer and Roland Reichle from the Distributed Systems Group, University of Kassel, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Communication and Collaboration in Heterogeneous Teams of Soccer Robots. In *Robotic Soccer*, chapter 1, pages 1–28. I-Tech Education and Publishing, Vienna, Austria, 2007. In collection [1552]. Online available at <http://s.i-techonline.com/Book/Robotic-Soccer/ISBN978-3-902613-21-9-rs01.pdf> [accessed 2008-04-23].

- [1564] Kurt Geihs. Why Robots Play Soccer (Keynote Extended Abstract). In Judith Bishop and Trish Alexander, editors, *Annual Conference of the South African Institute of Computer Scientists and Information Technologists (SAICSIT 2006)*, International Proceedings Series. SAICSIT, ACM, 2006. ISBN 1-59593-567-3.
- [1565] Philipp Andreas Baer, Roland Reichle, Michael Zapf, Thomas Weise, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. A Generative Approach to the Development of Autonomous Robot Software. In *Proceedings of 4th IEEE Workshop on Engineering of Autonomic and Autonomous Systems (EASE 2007)*. IEEE, March 26-29, 2007, Tucson, AZ U.S.A. ISBN 0-7695-2809-0. doi:10.1109/EASE.2007.2. Online available at <http://www.it-weise.de/documents/index.html#BRZWG2007AR> [accessed 2008-08-12].
- [1566] Ulrich Kaufmann, Roland Reichle, Christof Hoppe, and Philipp Andreas Baer. An Unsupervised Approach for Adaptive Color Segmentation. In *International Workshop on Robot Vision, VISAPP 2007*. Springer, 2007. Online available at <http://neuro.informatik.uni-ulm.de/basilic/Publications/2007/KRHB07/AdaptColorSeg-388.pdf> [accessed 2007-09-15].
- [1567] William J. Frawley, Gregory Piatetsky-Shapiro, and Christopher J. Matheus. Knowledge Discovery in Databases: An Overview. *AI Magazine*, pages 213–228, Fall 1992. ISSN 0738-4602. Online available at <http://citeseer.ist.psu.edu/524488.html> and <http://www.kdnuggets.com/gpspubs/aimag-kdd-overview-1992.pdf> [accessed 2007-08-11].
- [1568] David J. Hand, Heikki Mannila, and Padhraic Smyth. *Principles of Data Mining*. MIT Press, Cambridge, MA, August 2001. ISBN 0-262-08290-X.
- [1569] Joseph P. Bigus. *Data Mining With Neural Networks: Solving Business Problems from Application Development to Decision Support*. McGraw-Hill (Tx), May 20, 1996. ISBN 978-0070057791.
- [1570] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, USA, November 3, 1995. ISBN 978-0198538646.
- [1571] Vladimir N. Vapnik. *The Nature of Statistical Learning Theory*. Information Science and Statistics. Springer-Verlag, second edition, November 19, 1999 (first edition: 1995). ISBN 978-0387987804.
- [1572] Lipo Wang, editor. *Support Vector Machines: Theory and Applications*, volume 177 of *Studies in Fuzziness and Soft Computing*. Springer, August 2005. ISBN 978-3540243885. doi:10.1007/b95439.
- [1573] Christopher J. C. Burges from the Bell Laboratories, Lucent Technologies. A Tutorial on Support Vector Machines for Pattern Recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, 1998. Online available at <http://citeseer.ist.psu.edu/burges98tutorial.html> and <http://research.microsoft.com/~cburges/papers/SVMTutorial.pdf> [accessed 2007-08-08].
- [1574] Christiaan M. van der Walt and Etienne Barnard. Data characteristics that determine classifier performance. In *Proceedings of the Sixteenth Annual Symposium of the Pattern Recognition Association of South Africa*, pages 160–165, November 23–25, 2005. Online available at <http://www.meraka.org.za/pubs/CvdWalt.pdf> [accessed 2007-08-08].
- [1575] Alan Agresti. *An Introduction to Categorical Data Analysis*. Wiley-Interscience, first edition, January 1996. ISBN 978-0471113386.
- [1576] Michael J. A. Berry and Gordon S. Linoff. *Mastering Data Mining: The Art and Science of Customer Relationship Management*. Wiley, first edition, December 1999. ISBN 978-0471331230.
- [1577] Ian H. Witten and Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. The Morgan Kaufmann Series in Data Management Systems. Morgan Kaufmann, first edition, October 1999. ISBN 978-1558605527.

- [1578] Pedro Domingos and Michael Pazzani. On the Optimality of the Simple Bayesian Classifier under Zero-One Loss. *Machine Learning*, 29(2-3):103–130, November 1997. ISSN 0885-6125 (Print) 1573-0565 (Online). Online available at <http://citeseer.ist.psu.edu/domingos97optimality.html> and <http://www.ics.uci.edu/~pazzani/Publications/mlj97-pedro.pdf> [accessed 2007-08-11].
- [1579] Irina Rish. An empirical study of the naive Bayes classifier. In *Proceedings of IJCAI-01 workshop on Empirical Methods in AI, International Joint Conference on Artificial Intelligence*, pages 41–46. American Association for Artificial Intelligence, 2001. In proceedings [1953]. Online available at <http://www.cc.gatech.edu/~isbell/classes/reading/papers/Rish.pdf> [accessed 2007-08-11].
- [1580] Ran Avnimelech and Nathan Intrator. Boosting Regression Estimators. *Neural Computation*, 11(2):499–520, 1999. Online available at <http://neco.mitpress.org/cgi/reprint/11/2/499.pdf> and <http://citeseer.ist.psu.edu/avnimelech99boosting.html> [accessed 2007-9-15].
- [1581] Robert E. Schapire. The Strength of Weak Learnability. *Machine Learning*, 5:197–227, June 1990. ISSN 0885-6125 (Print) 1573-0565 (Online). doi:10.1007/BF00116037. Online available at <http://www.springerlink.com/content/x02406w7q5038735/fulltext.pdf> and <http://citeseer.ist.psu.edu/schapire90strength.html> [accessed 2007-09-15].
- [1582] Yoav Freund and Robert E. Schapire from the AT&T Labs, 180 Park Avenue, Florham Park, New Jersey 07932, USA. A decision-theoretic generalization of on-line learning and an application to boosting. In *European Conference on Computational Learning Theory*, volume 904 of *Lecture Notes In Computer Science (LNCS)*, pages 23–37. Springer, 1995. ISBN 3-540-59119-2. Online available at <http://citeseer.ist.psu.edu/freund95decisiontheoretic.html> and <http://www.cs.princeton.edu/~schapire/uncompress-papers.cgi/FreundSc95.ps> [accessed 2007-09-15]. See also [1954].
- [1583] Yoav Freund and Robert E. Schapire. Experiments with a New Boosting Algorithm. In *Proceedings of the Thirteenth International Conference on Machine Learning, ICML '96*, pages 148–156, July 3–6, 1996, Bari, Italy. Online available at <http://citeseerx.ist.psu.edu/viewdoc/summary?cid=45368> and <http://www.cs.ucsb.edu/~yfwang/courses/cs290i/papers/freund-adaboost96.pdf> [accessed 2008-06-23].
- [1584] Yoav Freund from the AT&T Labs, 180 Park Avenue, Florham Park, New Jersey 07932, USA. Boosting a Weak Learning Algorithm by Majority. In *COLT: Proceedings of the third annual Workshop on Computational Learning Theory*. Morgan Kaufmann Publishers, August 8–9, 1990, University of Rochester, Rochester, New York, USA. ISBN 1-55860-146-5. Online available at <http://citeseer.ist.psu.edu/freund95boosting.html> [accessed 2007-09-15].
- [1585] Robert E. Schapire, Yoav Freund, Peter Bartlett, and Wee Sun Lee. Boosting the margin: a new explanation for the effectiveness of voting methods. In *Proceedings 14th International Conference on Machine Learning*, pages 322–330. Morgan Kaufmann, July 1997, Nashville, Tennessee, USA. ISBN 1-55860-486-3. Online available at <http://citeseer.ist.psu.edu/schapire97boosting.html> [accessed 2007-09-15]. See also [1955].
- [1586] Carlo Ghezzi, Mehdi Jazayeri, and Dino Mandrioli. *Fundamentals of Software Engineering*. Pearson Education, Prentice Hall, second edition, September 19, 2002. ISBN 9780133056990.
- [1587] Jabir and J. W. Moore from the 7803 Whiterim Terrace, Potomac, MD 20854, USA. A search for fundamental principles of software engineering. *Computer Standards & Interfaces*, 19(2):155–160, March 1998. doi:10.1016/S0920-5489(98)00009-9. Online available at [http://dx.doi.org/10.1016/S0920-5489\(98\)00009-9](http://dx.doi.org/10.1016/S0920-5489(98)00009-9) and <http://www.gelog.etsmtl.ca/publications/pdf/249.pdf> [accessed 2007-09-02].

- [1588] Ingrid Wetzel from the Hamburg University, Computer Science Department, Vogt-Koelln-Straße 30, 22527 Hamburg, Germany. Information Systems Development with Anticipation of Change: Focussing on Professional Bureaucracies. In *Proceedings of Hawaii International Conference on Systems Sciences, HICSS 34*. IEEE Computer Society, January 2001, Maui, Hawaii, USA. Online available at <http://citeseer.ist.psu.edu/532081.html> and <http://swt-www.informatik.uni-hamburg.de/publications/download.php?id=177> [accessed 2007-09-02].
- [1589] Eric A. Marks and Michael Bell. *Executive's Guide to Service oriented architecture (SOA): A Planning and Implementation Guide for Business and Technology*. John Wiley & Sons, Inc., Hoboken, NJ, April 2006. ISBN 978-0-470-03614-3.
- [1590] Thomas Erl. *Service-Oriented Architecture (SOA): Concepts, Technology, and Design*. The Prentice Hall Service-Oriented Computing Series from Thomas Erl. Prentice Hall PTR, August 2, 2005. ISBN 978-0131858589.
- [1591] David Booth and Canyang Kevin Liu. *Web Services Description Language (WSDL) Version 2.0 Part 0: Primer*. World Wide Web Consortium (W3C), June 26, 2007. W3C Recommendation. Online available at <http://www.w3.org/TR/2007/REC-wsdl20-primer-20070626> [accessed 2007-09-02].
- [1592] Anupriya Ankolekar, Mark Burstein, Grit Denker, Daniel Elenius, Jerry Hobbs, Lalana Kagal, Ora Lassila, Drew McDermott, Deborah McGuinness, Sheila McIlraith, Massimo Paolucci, Bijan Parsia, Terry Payne, Marta Sabou, Craig Schlenoff, Evren Sirin, Monika Solanki, Naveen Srinivasan, Katia Sycara, and Randy Washington. *OWL-S 1.1 Release, OWL-based Web Service Ontology*. Web-Ontology Working Group at the World Wide Web Consortium, 2004. Online available at <http://www.daml.org/services/owl-s/1.1/> [accessed 2007-09-02].
- [1593] Dumitru Roman, Uwe Keller, and Holger Lausen. *WSMO – Web Service Modeling Ontology*. Digital Enterprise Research Institute (DERI), February 2004. Online available at <http://www.wsmo.org/2004/d2/v0.1/20040214/> [accessed 2007-09-02]. See also <http://www.wsmo.org/> [accessed 2007-09-02] and [1594].
- [1594] Dumitru Roman, Uwe Keller, Holger Lausen, Jos de Bruijn, Ruben Lara, Michael Stollberg, Axel Polleres, Cristina Feier, Christoph Bussler, and Dieter Fensel. Web Service Modeling Ontology. *Applied Ontology*, 1(1):77–106, 2005. See also <http://www.wsmo.org/> [accessed 2007-09-02] and [1593].
- [1595] Steffen Bleul and Thomas Weise from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. An Ontology for Quality-Aware Service Discovery. In C. Zirpins, G. Ortiz, W. Lamerdorf, and W. Emmerich, editors, *Engineering Service Compositions: First International Workshop, WESC05*, volume RC23821 of *IBM Research Report*. Yorktown Heights: IBM Research Division, December 12, 2005, Vrije Universiteit Amsterdam, The Netherlands. See also [1596]. Online available at <http://www.it-weise.de/documents/files/BW2005QASD.pdf> [accessed 2008-08-12].
- [1596] Steffen Bleul, Thomas Weise, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. An Ontology for Quality-Aware Service Discovery. *Computer Systems Science Engineering*, 21(4)(Special issue on “Engineering Design and Composition of Service-Oriented Applications”), July 2006. See also [1595]. Online available at <http://www.it-weise.de/documents/files/BWG2006QASD.pdf> [accessed 2008-08-12].
- [1597] Steffen Bleul and Kurt Geihs. ADDO: Automatic Service Brokering in Service Oriented Architectures, Project Homepage, 2007. Online available at <http://www.vs.uni-kassel.de/ADDO/> [accessed 2007-09-02].
- [1598] *METEOR-S: Semantic Web Services and Processes*. LSDIS Lab (Large Scale Distributed Information Systems), Department of Computer Science, University of Georgia, 2004. Online available at <http://lsdis.cs.uga.edu/projects/meteor-s/> [accessed 2007-09-02].

- [1599] Asunción Gómez-Pérez, Rafael González-Cabero, and Manuel Lama. ODE SWS: A Framework for Designing and Composing Semantic Web Services. *IEEE Intelligent Systems*, 19(4):24–31, July–August 2004. ISSN 1541-1672. doi:10.1109/MIS.2004.32. Online available at <http://iswc2004.semanticweb.org/demos/14/paper.pdf> [accessed 2007-09-02]. See also [1600].
- [1600] Asunción Gómez-Pérez, Rafael González-Cabero, and Manuel Lama. A Framework for Designing and Composing Semantic Web Services. In *Semantic Web Services, First International Semantic Web Services Symposium, Proceedings of 2004 AAAI Spring Symposium Series*, March 22–24, 2004, History Corner, main quad (Building 200), Stanford University, CA, USA. Online available at <http://www.daml.ecs.soton.ac.uk/SSS-SWS04/44.pdf> [accessed 2007-09-02]. See also [1599].
- [1601] M. Brian Blake, Kwok Ching Tsui, and Andreas Wombacher. The EEE-05 challenge: a new Web service discovery and composition competition. In *Proceedings of the 2005 IEEE International Conference on e-Technology, e-Commerce, and e-Service, EEE'05*, pages 780–783, March 29–April 1, 2005. doi:10.1109/EEE.2005.131. Online available at <http://ws-challenge.georgetown.edu/ws-challenge/The%20EEE.htm> [accessed 2007-09-02].
- [1602] M. Brian Blake, William K.W. Cheung, Michael C. Jaeger, and Andreas Wombacher. WSC-06: The Web Service Challenge. In *Proceedings of 2006 IEEE Joint Conference on E-Commerce Technology (CEC'06) and Enterprise Computing, E-Commerce and E-Services (EEE'06)*, pages 505–508, 2006. In proceedings [1956].
- [1603] M. Brian Blake, William K.W. Cheung, Michael C. Jaeger, and Andreas Wombacher. WSC-07: Evolving the Web Service Challenge. In *Proceedings of IEEE Joint Conference (CEC/EEE 2007) on E-Commerce Technology (9th CEC'07) and Enterprise Computing, E-Commerce and E-Services (4th EEE'07)*, pages 422–423, 2006. In proceedings [1957].
- [1604] Steffen Bleul, Thomas Weise, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Large-Scale Service Composition in Semantic Service Discovery. In Ws-Challenge Part: M. Brian Blake, Andreas Wombacher, Michel C. Jaeger, and William K. Cheung, editors, *Proceedings of 2006 IEEE Joint Conference on E-Commerce Technology (CEC'06) and Enterprise Computing, E-Commerce and E-Services (EEE'06)*, pages 427–429, 2006. In proceedings [1956]. 1st place in 2006 WSC. Online available at <http://www.it-weise.de/documents/files/BWG2006WSC.pdf> [accessed 2008-08-12]. See 2007 WSC [1605] and [1606].
- [1605] Steffen Bleul, Thomas Weise, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Making a Fast Semantic Service Composition System Faster. In *Proceedings of IEEE Joint Conference (CEC/EEE 2007) on E-Commerce Technology (9th CEC'07) and Enterprise Computing, E-Commerce and E-Services (4th EEE'07)*, pages 517–520, 2007. In proceedings [1957]. 2nd place in 2007 WSC. Online available at <http://www.it-weise.de/documents/files/BWG2007WSC.pdf> [accessed 2008-08-12]. See 2006 WSC [1604] and [1606].
- [1606] Thomas Weise, Steffen Bleul, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Web Service Composition Systems for the Web Service Challenge – A Detailed Review. *Kasseler Informatikschriften (KIS)* 2007, 7, University of Kassel, FB16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany, University of Kassel, November 19, 2007. Persistent Identifier: urn:nbn:de:hebis:34-2007111919638. Online available at <http://kobra.bibliothek.uni-kassel.de/handle/urn:nbn:de:hebis:34-2007111919638> and <http://www.it-weise.de/documents/files/WBG2007WSCb.pdf> [accessed 2007-11-20]. See also [1604, 1605].

- [1607] Thomas Weise, Steffen Bleul, Diana Comes, and Kurt Geihs from the Distributed Systems Group, FB 16 – Elektrotechnik und Informatik, University of Kassel, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Different Approaches to Semantic Web Service Composition. In Abdelhamid Mellouk, Jun Bi, Guadalupe Ortiz, Dickson K. W. Chiu, and Manuela Popescu, editors, *Proceedings of The Third International Conference on Internet and Web Applications and Services, ICIW 2008*, pages 90–96, June 8–13, 2008, Athens, Greece. IEEE Computer Society Press, Los Alamitos, CA, USA. ISBN 978-0-7695-3163-2. Product Number: E3163. BMS Part Number: CFP0816C-CDR. Library of Congress Number: 2008922600. Online available at <http://www.it-weise.de/documents/index.html#WBCG2008ICIW> [accessed 2008-08-12].
- [1608] David C. Fallside and Priscilla Walmsley. *XML Schema Part 0: Primer Second Edition*. World Wide Web Consortium (W3C), second edition, October 28, 2004. W3C Recommendation. Online available at <http://www.w3.org/TR/2004/REC-xmlschema-0-20041028/> [accessed 2007-09-02].
- [1609] Tim Bray, Jean Paoli, C. M. Sperberg-McQueen, Eve Maler, and François Yergeau. *Extensible Markup Language (XML) 1.0 (Fourth Edition)*. World Wide Web Consortium (W3C), September 29, 2007. W3C Recommendation. Online available at <http://www.w3.org/TR/2006/REC-xml-20060816> [accessed 2007-09-02].
- [1610] *BPEL4WS, Business Process Execution Language for Web Services Version 1.1*. IBM, BEA Systems, Microsoft, SAP AG, Siebel Systems, May 2003. Online available at <http://www.ibm.com/developerworks/library/specification/ws-bpel/> [accessed 2007-09-03].
- [1611] Diane Jordan, John Evdemon, Alexandre Alves, Assaf Arkin, Sid Askary, Charlton Barreto, Ben Bloch, Francisco Curbera, Mark Ford, Yaron Goland, Alejandro Guízar, Neelakantan Kartha, Canyang Kevin Liu, Rania Khalaf, Dieter König, Mike Marin, Vinkesh Mehta, Satish Thatte, Danny van der Rijn, Prasad Yendluri, and Alex Yiu. *Web Services Business Process Execution Language Version 2.0*. Organization for the Advancement of Structured Information Standards (OASIS), April 11, 2007. Online available at <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0.pdf> [accessed 2007-10-25]. Technical Committee: OASIS Web Services Business Process Execution Language (WSBPEL) TC.
- [1612] Thomas Weise from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. *Global Optimization Algorithms – Theory and Application*. Thomas Weise, second edition, 2008. Online available at <http://www.it-weise.de/> [accessed 2008-08-12].
- [1613] D. G. Kleinbaum, L. L. Kupper, and K. E. Muller, editors. *Applied regression analysis and other multivariable methods*. PWS Publishing Co., Boston, MA, USA, 1988. ISBN 0-871-50123-6.
- [1614] John Fox. *Applied Regression Analysis, Linear Models, and Related Methods*. Sage Publications, Inc, New York, subsequent edition, February 1997. ISBN 978-0803945401.
- [1615] D. M. Ellis. Book Reviews: “Applied Regression Analysis” by N. P. Draper and H. Smith. *Applied Statistics*, 17(1):83–84, 1968. ISSN 0035-9254. Reviews [1616].
- [1616] Norman Richard Draper and H. Smith. *Applied regression analysis*. Wiley, New York, 1966. Reviewed in [1615].
- [1617] Abdellah Salhi, Hugh Glaser, and David De Roure. Parallel Implementation of a Genetic-Programming based Tool for Symbolic Regression. Technical Report DSSE Technical Reports: DSSE-TR-97-3, Declarative Systems & Software Engineering Group, Department of Electronics and Computer Science, University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom, July 18, 1997. See also [1618]. Online available at <http://www.dsse.ecs.soton.ac.uk/techreports/95-03/97-3.html> [accessed 2007-09-09].

- [1618] Abdellah Salhi, Hugh Glaser, and David De Roure. Parallel Implementation of a Genetic-Programming Based Tool for Symbolic Regression. *Information Processing Letters*, 66(6):299–307, June 30, 1998. ISSN 0020-0190. Online available at [http://dx.doi.org/10.1016/S0020-0190\(98\)00056-8](http://dx.doi.org/10.1016/S0020-0190(98)00056-8) [accessed 2007-09-09]. See also [1617].
- [1619] John Duffy and Jim Engle-Warnick. Using Symbolic Regression to Infer Strategies from Experimental Data. In David A. Belsley and Christopher F. Baum, editors, *Fifth International Conference: Computing in Economics and Finance*, page 150, June 24–26, 1999. Boston College, MA, USA. Online available at <http://www.pitt.edu/~jduffy/papers/Usr.pdf> and <http://citeseer.ist.psu.edu/304022.html> [accessed 2007-09-09]. Later published as bookchapter: [1620].
- [1620] John Duffy and Jim Engle-Warnick. Using Symbolic Regression to Infer Strategies from Experimental Data. In *Evolutionary Computation in Economics and Finance*, chapter 4, pages 61–84. Physica-Verlag Heidelberg, 2002. In collection [458]. Presented at CEF’99 (see [1619]). Online available at <http://www.pitt.edu/~jduffy/docs/Usr.ps> [accessed 2007-09-09].
- [1621] Maarten Keijzer from the Zevenwouden 186, 3524CX Utrecht, The Netherlands. Scaled Symbolic Regression. *Genetic Programming and Evolvable Machines*, 5(3):259–269, September 2004. ISSN 1389-2576 (Print) 1573-7632 (Online). doi:10.1023/B:GENP.0000030195.77571.f9. Online available at <http://www.springerlink.com/content/x035121165125175/fulltext.pdf> and <http://dx.doi.org/10.1023/B:GENP.0000030195.77571.f9> [accessed 2007-09-09].
- [1622] Zbigniew Michalewicz. A Step Towards Optimal Topology of Communication Networks. In Vibeke Libby, editor, *Proceedings of Data Structures and Target Classification, the SPIE’s International Symposium on Optical Engineering and Photonics in Aerospace Sensing*, volume 1470 of *SPIE*, pages 112–122. SPIE – The International Society for Optical Engineering., April 1–5 1991, Orlando, Florida. ISBN 0819405795, 978-0819405791.
- [1623] Samuel Pierre. Inferring new design rules by machine learning: a case study of topological optimization. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 28(5):575–585, 1998.
- [1624] Susan Coombs and Lawrence Davis. Genetic algorithms and communication link speed design: constraints and operators. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 257–260, 1987. In proceedings [750].
- [1625] Aaron Kershenbaum. *Telecommunications Network Design Algorithms*. McGraw-Hill Companies, January 15, 1993. ISBN 0070342288, 978-0070342286.
- [1626] Faris N. Abuali, Dale A. Schoenefeld, and Roger L. Wainwright. Terminal assignment in a communications network using genetic algorithms. In *CSC’94: Proceedings of the 22nd annual ACM computer science conference on Scaling up : meeting the challenge of complexity in real-world computing applications*, pages 74–81, 1994, Phoenix, Arizona, United States. ACM, New York, NY, USA. ISBN 0-89791-634-4. doi:10.1145/197530.197559.
- [1627] Katharina Anna Zweig from the Fakultät für Informations- und Kognitionswissenschaften der Eberhard-Karls-Universität Tübingen. *On Local Behavior and Global Structures in the Evolution of Complex Networks*. PhD thesis, Fakultät für Informations- und Kognitionswissenschaften der Eberhard-Karls-Universität Tübingen, July 2007. Online available at <http://www-pr.informatik.uni-tuebingen.de/mitarbeiter/katharinazweig/downloads/Diss.pdf> [accessed 2008-06-12]. Zweig was formerly known as Lehmann.
- [1628] General Secretariat of the Council of the European Union; COST Secretariat. *COST 1991–1992*, volume 7 of *Collected Agreements Concluded within the Framework of European Cooperation in the Field of Scientific and Technical Research*. Office for

- Official Publications of the European Communities, Luxembourg, 1994. ISBN 92-824-1119-2, 92-824-1120-6.
- [1629] M. O'Mahony, Mark C. Sinclair, and B. Mikac. Ultra-High Capacity Optical Transmission Network: European Research Project COST 239. *Information, Telecommunications, Automata Journal*, 12(1-3):33-45, 1993.
 - [1630] Lawrence Davis Louis Anthony Cox, Jr. and Yuping Qiu. Dynamic Anticipatory Routing In Circuit-Switched Telecommunications Networks. In *Handbook of Genetic Algorithms*, pages 124-143. Van Nostrand Reinhold, 1991. In collection [752].
 - [1631] *The gnutella protocol specification v0.4*. Clip2 Distributed Search Services, August 10, 2000. Online available at http://www.stanford.edu/class/cs244b/gnutella_protocol_0.4.pdf [accessed 2008-06-24]. Document Revision 1.2, June 15, 2001 online available at http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf [accessed 2008-06-24].
 - [1632] Matei Ripeanu, Ian Foster, and Adriana Iamnitchi. Mapping the gnutella network: Properties of large-scale peer-to-peer systems and implications for system design. *IEEE Internet Computing Journal*, 6(1):50-57, February 2002. Online available at <http://people.cs.uchicago.edu/~matei/papers/ic.pdf> [accessed 2008-06-24]. <http://citeseer.ist.psu.edu/ripeanu02mapping.html>
 - [1633] Edsger Wybe Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269-271, 1959. Online available at <http://www-m3.ma.tum.de/twiki/pub/MN0506/WebHome/dijkstra.pdf> [accessed 2008-06-24].
 - [1634] C. Hedrick from the Rutgers University. Routing Information Protocol. Request for Comments (RFC) 1058, Internet Engineering Task Force (IETF), Internet Society (ISOC), Network Working Group, June 1988. Online available at <http://tools.ietf.org/html/rfc1058> and <http://www.faqs.org/rfcs/rfc1058.html> [accessed 2008-06-27].
 - [1635] John Moy from the Proteon, Inc., 2 Technology Drive, Westborough, MA 01581-5008. The OSPF Specification. Request for Comments (RFC) 1131, Internet Engineering Task Force (IETF), Internet Society (ISOC), October 1989. Online available at <http://www.packetizer.com/rfc/rfc1131/pdf> [accessed 2008-06-27].
 - [1636] Masaharu Munetomo, Yoshiaki Takai, and Yoshiharu Sato. An Adaptive Network Routing Algorithm Employing Path Genetic Operators. In *Proceedings of the Seventh International Conference on Genetic Algorithms*, pages 643-649, 1997. In proceedings [746].
 - [1637] Masaharu Munetomo, Yoshiaki Takai, and Yoshiharu Sato. An adaptive routing algorithm with load balancing by a genetic algorithm. *Transactions of the Information Processing Society of Japan (IPSJ)*, 39(2):219-227, 1998. in Japanese.
 - [1638] Masaharu Munetomo, Yoshiaki Takai, and Yoshiharu Sato. A Migration Scheme for the Genetic Adaptive Routing Algorithm. In *Proceedings of the 1998 IEEE Conference on Systems, Man, and Cybernetics*, volume 3, pages 2774-2779. IEEE Press, October 11-14, 1998, San Diego, U.S.A. doi:10.1109/ICSMC.1998.725081.
 - [1639] Masaharu Munetomo from the Graduate School of Engineering, Hokkaido University, North 13, West 8, Sapporo 060-8628, Jappan. Designing genetic algorithms for adaptive routing algorithms in the internet. In *Evolutionary Telecommunications: Past, Present, and Future*, pages 215-216, 1999. Online available at <http://citeseer.ist.psu.edu/198441.html> and http://uk.geocities.com/markcsinclair/ps/etppf_mun.ps.gz [accessed 2008-06-25]. Presentation online available at http://uk.geocities.com/markcsinclair/ps/etppf_mun_slides.ps.gz and <http://assam.cims.hokudai.ac.jp/~munetomo/documents/gecco99-ws.ppt> [accessed 2008-06-25].
 - [1640] Ruud Schoonderwoerd, Owen Holland, and Janet Bruten. Ant-like agents for load balancing in telecommunications networks. In Jörg Mller, editor, *AGENTS'97: Proceedings of the first international conference on Autonomous agents*, pages 209-216,

- February 5–8 1997, Marina del Rey, California, United States. ACM, New York, NY, USA. ISBN 0-89791-877-0. doi:10.1145/267658.267718. Online available at <http://staff.washington.edu/paymana/swarm/schoon97-icaa.pdf> and <http://doi.acm.org/10.1145/267658.267718> [accessed 2008-07-26]. Slides online available at [accessed 2008-07-26]. See also <http://sigart.acm.org/proceedings/agents97/> [accessed 2008-07-26].
- [1641] Ralf Klasing, Burkhard Monien, R. Peine, and Elena Stöhr. Broadcasting in butterfly and de Bruijn networks. In Alain Finkel and Matthias Jantzen, editors, *Proceedings of 9th Annual Symposium on Theoretical Aspects of Computer Science (STACS'92)*, volume 577 of *Lecture Notes in Computer Science (LNCS)*, pages 351–362, February 13–15, 1992, Cachan, France. Springer, Berlin, Heidelberg, New York. ISBN 978-3-540-55210-9.
- [1642] Jiri Jaros and Vaclav Dvorak from the Brno University of Technology, Bozotechnova 2, 612 66, Brno, CZ. An Evolutionary Design Technique for Collective Communications on Optimal Diameter-Degree Networks. In *Genetic and Evolutionary Computation Conference (GECCO)*, 2008. In proceedings [550].
- [1643] Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, California 90291, USA. Internet Protocol, Darpa Internet Program Protocol Specification. RFC 791, Defense Advanced Research Projects Agency, Information Processing Techniques Office, 1400 Wilson Boulevard, Arlington, Virginia 22209, USA, September 1981. Online available at <http://tools.ietf.org/html/rfc791> [accessed 2008-06-13].
- [1644] Vinton Cerf, Yogen Dalal, and Carl Sunshine. Specification of Internet Transmission Control Program (December 1974 Version). Request for Comments (RFC) 675, Network Working Group, December 1974. Online available at <http://tools.ietf.org/html/rfc675> [accessed 2008-06-13].
- [1645] Hirozumi Yamaguchi, Kozo Okano, Teruo Higashino, and Kenichi Taniguchi from the Department of Information and Computer Sciences, Osaka University, Machikaneyama 1-3, Toyonaka, Osaka 560, Japan. Synthesis of Protocol Entities' Specifications from Service Specifications in a Petri Net Model with Registers. In *ICDCS'95: Proceedings of the 15th International Conference on Distributed Computing Systems*, pages 510–517, May 30–June 2, 1995, Vancouver, British Columbia, Canada. IEEE Computer Society, Washington, DC, USA. ISBN 0-8186-7025-8. Online available at <http://citeseer.ist.psu.edu/yamaguchi95synthesis.html> [accessed 2008-06-15].
- [1646] Sérgio Granato de Araújo, Aloysio de Castro Pinto Pedroza, and Antônio Carneiro de Mesquita Filho from the Electrical Engineering Department, Universidade Federal do Rio de Janeiro (UFRJ), Caixa Postal 68504, CEP 21945-970, Rio de Janeiro, RJ, Brazil. Uma Metodologia de Projeto de Protocolos de Comunicação Baseada em Técnicas Evolutivas. In *XX Simpósio Brasileiro de Telecomunicações (SBT'03)*, October 5–8, 2003, Rio de Janeiro, Brasil. Online available at <http://www.gta.ufrj.br/ftp/gta/TechReports/AMP03e.pdf> [accessed 2008-06-21]. See also [1647] and the English version [859].
- [1647] Sérgio Granato de Araújo, Antônio Carneiro de Mesquita Filho, and Aloysio de Castro Pinto Pedroza from the Electrical Engineering Department, Universidade Federal do Rio de Janeiro (UFRJ), Caixa Postal 68504, CEP 21945-970, Rio de Janeiro, RJ, Brazil. Síntese de Circuitos Digitais Otimizados via Programação Genética. In *XXX Seminário Integrado de Software e Hardware (SEMISH'03) in Proceedings of XXIII Congresso da Sociedade Brasileira de Computação*, volume III, pages 273–285, August 4–5, 2003, Campinas, SP, Brasil. Online available at <http://www.gta.ufrj.br/ftp/gta/TechReports/AMP03d.pdf> [accessed 2008-06-21]. See also [1646] and the English version [859].

- [1648] Nicholas Peter Sharples. Evolutionary Approaches to Adaptive Protocol Design. In Darren Pearce, editor, *CSRP 512: The 12th White House Papers Graduate Research in Cognitive and Computing Sciences at Sussex*, Cognitive Science Research Papers, pages 60–62. University of Sussex at Brighton, November 1999. Online available at <ftp://ftp.informatics.sussex.ac.uk/pub/reports/csrp/csrp512.ps.Z> [accessed 2008-06-23].
- [1649] Christian F. Tschudin and Lidia A. R. Yamamoto. A Metabolic Approach to Protocol Resilience. In *Autonomic Communication – Revised Selected Papers from the First International IFIP Workshop, WAC 2004*, volume 3457/2005 of *Lecture Notes in Computer Science (LNCS)*, pages 191–206. Springer Berlin / Heidelberg, October 18–19, 2004, Berlin, Germany. ISBN 978-3-540-27417-9. doi:10.1007/11520184_15. Online available at <http://cn.cs.unibas.ch/pub/doc/2004-wac-lncs.pdf> [accessed 2008-06-20], presentation online available at <http://www.autonomic-communication.org/wac/wac2004/program/presentations/wac2004-tschudin-yamamoto-slides.pdf> [accessed 2008-06-20].
- [1650] Lidia A. R. Yamamoto and Christian F. Tschudin from the Computer Science Department, University of Basel, Bernoullistrasse 16, CH-4056 Basel, Switzerland. Experiments on the Automatic Evolution of Protocols Using Genetic Programming. In *Autonomic Communication – Revised Selected Papers from the Second International IFIP Workshop, WAC 2005*, volume 3854/2006 of *Lecture Notes in Computer Science (LNCS)*, pages 13–28. Springer Berlin / Heidelberg, October 2–5 2005, Vouliagmeni, Athens, Greece. ISBN 978-3-540-32992-3. doi:10.1007/11687818_2. Online available at <http://cn.cs.unibas.ch/people/ly/doc/wac2005-lyct.pdf> [accessed 2008-06-20]. See also [1651].
- [1651] Lidia A. R. Yamamoto and Christian F. Tschudin from the Computer Science Department, University of Basel, Bernoullistrasse 16, CH-4056 Basel, Switzerland. Experiments on the Automatic Evolution of Protocols Using Genetic Programming. Technical Report CS-2005-002, University of Basel, April 21, 2005. Online available at <http://cn.cs.unibas.ch/people/ly/doc/wac2005tr-lyct.pdf> [accessed 2008-06-20]. See also [1650].
- [1652] Daniele Miorandi, Paolo Dini, Eitan Altman, and Hisao Kameda. *WP 2.2 – Paradigm Applications and Mapping, D2.2.2 Framework for Distributed On-line Evolution of Protocols and Services*. BIOlogically inspired NETwork and Services (BIONETS) and Future and Emerging Technologies (FET) project of the EU, 2nd edition, June 14, 2007. BIONETS/CN/wp2.2/v2.1. Editor: Lidia A. R. Yamamoto. Verification David Linner and Fran coise Baude. Status: Final, Public. Online available at http://www.bionets.eu/docs/BIONETS_D2_2_2.pdf and 2008-06-23 [accessed .]
- [1653] Richard Heady, George Luger, Arthur Maccabe, and Mark Servilla. The architecture of a network level intrusion detection system. Technical Report CS90-20, Department of Computer Science, University of New Mexico, USA, August 1990.
- [1654] Gianluigi Folino, Clara Pizzuti, and Giandomenico Spezzano from the ICAR-CNR, Via P.Bucci 41/C, Univ. della Calabria, 87036 Rende (CS), Italy. GP Ensemble for Distributed Intrusion Detection Systems. In *Pattern Recognition and Data Mining, Proceedings of the Third International Conference on Advances in Pattern Recognition, ICAPR 2005*, volume 3686/2005 of *Lecture Notes in Computer Science (LNCS)*, pages 54–62. Springer Berlin / Heidelberg, August 22–25, 2005, Bath, UK. ISBN 978-3-540-28757-5. doi:10.1007/11551188_6. Online available at <http://dns2.icar.cnr.it/pizzuti/icapr05.pdf> [accessed 2008-06-23].
- [1655] Srinivas Mukkamala, Andrew H. Sung, and Ajith Abraham. Modeling intrusion detection systems using linear genetic programming approach. In Robert Orchard, Chunsheng Yang, and Moonis Ali, editors, *Proceedings of the 17th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, IEA/AIE 2004*, volume 3029/2004 of *Lecture Notes in Computer*

- Science (LNCS)*, pages 633–642. Springer Springer Verlag Inc, May 17–20, 2004, Ottawa, Canada. ISBN 3-540-22007-0, 978-3-540-22007-7. doi:10.1007/b97304. Online available at <http://www.rmltech.com/LGP%20Based%20IDS.pdf> [accessed 2008-06-17].
- [1656] James V. Hansen, Paul Benjamin Lowry, Rayman Meservy, and Dan McDonald. Genetic programming for prevention of cyberterrorism through previous dynamic and evolving intrusion detection. *Decision Support Systems*, 43(4):1362–1374, August 2007. doi:10.1016/j.dss.2006.04.004. Special Issue Clusters. Online available at <http://dx.doi.org/10.1016/j.dss.2006.04.004> and <http://ssrn.com/abstract=877981> [accessed 2008-06-17].
- [1657] Dawn Xiaodong Song from the Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA. Athena: an Automatic Checker for Security Protocol Analysis. In *Proceedings of 12th IEEE Computer Security Foundation Workshop*. IEEE press, June 1999. Online available at <http://www.ece.cmu.edu/~dawnsong/papers/Athena.pdf> and <http://www.ece.cmu.edu/~dawnsong/papers/Athena.ps> [accessed 2008-06-23].
- [1658] ISO. Information technology – Security techniques – Entity authentication – Part 3: Mechanisms using digital signature techniques. ISO/IEC 9798-3, International Standards Organization (ISO), 1993. JTC 1 Information technology. TC/SC: JTC 1/SC 27. Status: withdrawn.
- [1659] Tadashi Nakano and Tatsuya Suda. Applying biological principles to designs of network services. *Applied Soft Computing*, 7(3):870–878, June 2007. ISSN 1568-4946. doi:10.1016/j.asoc.2006.04.006.
- [1660] Paul Grace. Genetic Programming and Protocol Configuration. Master’s thesis, Computing Department, Lancaster University, September 2000. Supervisor: Professor Gordon Blair. Online available at <http://www.lancs.ac.uk/postgrad/gracep/msc.pdf> [accessed 2008-06-23].
- [1661] *JavaGroups – A Reliable Multicast Communication Toolkit for Java*, Cornell University and SourceForge, 1999–2008. See <http://www.cs.cornell.edu/Info/Projects/JavaGroupsNew/> [accessed 2008-06-23] and <http://sourceforge.net/projects/javagroups> [accessed 2008-06-23].
- [1662] Werner Van Belle, Tom Mens, and Theo DHondt from the Programming Technology Lab, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium. Using Genetic Programming to Generate Protocol Adaptors for Interprocess Communication. In Andy M. Tyrrell, Pauline C. Haddow, and Jim Torresen, editors, *Proceedings of the 5th International Conference on Evolvable Systems: From Biology to Hardware. ICES 2003*, volume 2606/2003 of *Lecture Notes in Computer Science (LNCS)*, pages 67–73. Springer Berlin / Heidelberg, March 17–20, 2003, Trondheim, Norway. ISBN 978-3-540-00730-2. Online available at <http://werner.yellowcouch.org/Papers/genadap03/ICES-online.pdf> and <http://prog.vub.ac.be/Publications/2003/vub-prog-tr-03-33.pdf> [accessed 2008-06-23].
- [1663] Werner Van Belle from the Programming Technology Lab (PROG), Department Informatica (DINF), Vrije Universiteit Brussel (VUB), Pleinlaan 2, 1050 Brussels, Belgium. *Automatic Generation of Concurrency Adaptors by Means of Learning Algorithms*. PhD thesis, Vrije Universiteit Brussel, Faculteit van de Wetenschappen, Departement Informatica, August 2001. Promotor: Theo D’Hondt, co-Promotor: Tom Tourwé. Online available at <http://borg.rave.org/adaptors/Thesis6.pdf> [accessed 2008-06-23].
- [1664] Theodore S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall PTR, 2001, (2nd ed: January 2002). ISBN 978-0130422323, 0130422320. Partly online available at <http://books.google.com/books?id=8ZXdHQAACAAJ> [accessed 2008-07-27].
- [1665] Andrea Goldsmith from the Stanford University, California. *Wireless Communications*. Cambridge University Press, The Edinburgh Building, Shaftesbury Road,

- Cambridge, CB2 8RU, UK, August 2005. Partly online available at http://books.google.com/books?id=mF7X_8D7_BIC [accessed 2008-07-27].
- [1666] Andreas F. Molisch. *Wireless Communications*. John Wiley & Sons / IEEE Press, November 18, 2005. ISBN 978-0470848876, 0470848871. Partly online available at <http://books.google.com/books?id=2NNCHgAACAAJ> [accessed 2008-07-27].
- [1667] Mohammad Ilyas and Imad Mahgoub, editors. *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*. CRC, July 16, 2004. ISBN 978-0849319686.
- [1668] Edgar H. Callaway, Jr. *Wireless Sensor Networks: Architectures and Protocols*. AUERBACH, August 26, 2003. ISBN 978-0849318238.
- [1669] Richard Harper, editor. *Inside the smart home*. Springer, 2003. ISBN 1852336889, 978-1-85233-688-2. Partly online available at <http://books.google.com/books?id=SvcHvHuv86gC> [accessed 2008-07-27].
- [1670] Werner Harke. *Smart Home – Vernetzung von Haustechnik und Kommunikationssystemen im Wohnungsbau*. Hüthig Verlag / C. F. Müller, December 31, 2003. ISBN 978-3-7880-7713-6, 3788077131.
- [1671] Adam Greenfield. *Everyware: The Dawning Age of Ubiquitous Computing*. New Riders Publishing, March 20, 2006. ISBN 0321384016, 978-0321384010.
- [1672] John Krumm, Gregory D. Abowd, Aruna Seneviratne, and Thomas Strang, editors. *UbiComp 2007: Ubiquitous Computing. Proceedings of the 9th International Conference*, volume 4717 of *Lecture Notes in Computer Science*, September 16–19, 2007, Innsbruck, Austria. Springer Berlin / Heidelberg. ISBN 978-3-540-74852-6. doi:10.1007/978-3-540-74853-3.
- [1673] *Workshop on Amorphous Computing*, September 13–14, 1999, DARPA ITO, MIT Artificial Intelligence Laboratory, Cambridge, MA, USA. MIT AI Lab. Online available at <http://www-swiss.ai.mit.edu/projects/amorphous/workshop-sept-99/> [accessed 2008-07-27].
- [1674] William J. Butera. *Programming a Paintable Computer*. PhD thesis, MIT Media Lab, Massachusetts Institute of Technology, MA, USA, February 2002. Online available at <http://www-swiss.ai.mit.edu/projects/amorphous/papers/butera-phd.pdf> [accessed 2008-07-27].
- [1675] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Computer Graphics*, 4(10):25–34, July 1987. ISSN 0097-8930. Online available at <http://graphics.ucmerced.edu/~mkallmann/courses/papers/Reynolds1987-flocks.pdf> and <http://www.cs.toronto.edu/~dt/siggraph97-course/cwr87/> [accessed 2008-07-27]. See also [1958].
- [1676] Leonid Libkin and Limsoon Wong. Query Languages for Bags and Aggregate Functions. *Journal of Computer and System Sciences*, 55(2):241–272, October 1997. Online available at <http://citeseer.ist.psu.edu/libkin97query.html> and <http://www.cs.toronto.edu/~libkin/papers/jcss97.ps.gz> [accessed 2007-09-12].
- [1677] Gultekin Özsoyoglu, Z. Meral Özsoyoglu, and Victor Matos. Extending Relational Algebra and Relational Calculus with Set-Valued Attributes and Aggregate Functions. *ACM Transactions on Database Systems*, 12(4):566–592, 1987. Online available at <http://doi.acm.org/10.1145/32204.32219> [accessed 2007-09-12].
- [1678] Anthony C. Klug. Equivalence of Relational Algebra and Relational Calculus Query Languages Having Aggregate Functions. *Journal of the ACM (JACM)*, 29(3):699–717, July 1982. ISSN 0004-5411. Online available at <http://doi.acm.org/10.1145/322326.322332> [accessed 2007-07-28].
- [1679] Ines Fernando Vega Lopez, Richard T. Snodgrass, and Bongki Moon. Spatiotemporal Aggregate Computation: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, 17(2):271–286, February 2005. ISSN 1041-4347. Online available at <http://citeseer.ist.psu.edu/634034.html> [accessed 2007-09-12]. <http://www.cs.arizona.edu/bkmoon/papers/tkde-stagg.pdf>

- [1680] Arbee L. P. Chen, Jui-Shang Chiu, and Frank S.C. Tseng. Evaluating Aggregate Operations Over Imprecise Data. *IEEE Transactions on Knowledge and Data Engineering*, 8(2):273–284, April 1996. ISSN 1041-4347. Online available at http://make.cs.nthu.edu.tw/alp/alp_paper/Evaluating%20aggregate%20operations%20over%20imprecise%20data.pdf [accessed 2007-09-12].
- [1681] Tina Dell’Armi, Wolfgang Faber, Giuseppe Ielpa, Nicola Leone, and Gerald Pfeifer. Aggregate Functions in Disjunctive Logic Programming Semantics, Complexity, and Implementation in DLV. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI) 2003*, pages 847–852, 2003. Online available at <http://citeseer.ist.psu.edu/643941.html> and <http://www.icons.rodan.pl/publications/%5BDellArmi2003%5D.pdf> [accessed 2007-09-12]. In proceedings [1959].
- [1682] John Miles Smith and Diane C. P. Smith. Database abstractions: Aggregation and Generalization. *ACM Transaction on Database Systems*, 2(2):105–133, June 1977. ISSN 0362-5915. Online available at <http://portal.acm.org/citation.cfm?doid=320544.320546> and <http://libra.msra.cn/paperDetail.aspx?id=794625&> [accessed 2007-08-11].
- [1683] Robbert van Renesse. The Importance of Aggregation. In A. Schiper, A. A. Shvartsman, H. Weatherspoon, and B. Y. Zhao, editors, *Future Directions in Distributed Computing*, volume 2584/2003 of *Lecture Notes in Computer Science (LNCS)*, pages 87–92. Springer Berlin/Heidelberg, 2003.
- [1684] Márk Jelasity, Alberto Montresor, and Özalp Babaoğlu. Gossip-based aggregation in large dynamic networks. *ACM Trans. Comput. Syst.*, 23(3):219–252, 2005. ISSN 0734-2071. Online available at <http://www.cs.unibo.it/bison/publications/aggregation-tocs.pdf> and <http://portal.acm.org/citation.cfm?id=1082470> [accessed 2007-08-01].
- [1685] David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-Based Computation of Aggregate Information. In *Proceedings of 44th Symposium on Foundations of Computer Science (FOCS 2003)*, pages 482–491, October 11–14, 2003, Cambridge, MA, USA. IEEE Computer Society, Los Alamitos, CA, USA. ISBN 0-7695-2040-5. Online available at <http://www.cs.cornell.edu/johannes/papers/2003/focs2003-gossip.pdf> and <http://citeseer.ist.psu.edu/kempe03gossipbased.html> [accessed 2007-09-15].
- [1686] Thomas Weise, Michael Zapf, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Evolving Proactive Aggregation Protocols. In *Genetic Programming – Proceedings of the 11th European Conference on Genetic Programming, EuroGP 2008*, pages 254–265, 2008. doi:10.1007/978-3-540-78671-9_22. In proceedings [901]. Online available at <http://www.it-weise.de/documents/index.html#WZG2008DGPFA> and http://dx.doi.org/10.1007/978-3-540-78671-9_22 [accessed 2008-08-12].
- [1687] Márk Jelasity and Alberto Montresor. Epidemic-Style Proactive Aggregation in Large Overlay Networks. In *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS’04)*, pages 102–109, March 2004. IEEE Computer Society, Tokyo, Japan. Online available at <http://citeseer.ist.psu.edu/jelasity04epidemicstyle.html> and <http://en.scientificcommons.org/9048443> [accessed 2007-08-13].
- [1688] Wendi Rabiner Heinzelman, Joanna Kulik, and Hari Balakrishnan. Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. In *MobiCom’99: Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pages 174–185, 1999, Seattle, Washington, United States. ACM Press, New York, NY, USA. ISBN 1-58113-142-9. Online available at <http://citeseer.ist.psu.edu/kulik99adaptive.html> and <http://www.comet.columbia.edu/~campbell/e6906/papers/heinzelman99.pdf> [accessed 2007-08-13].

- [1689] Matthew Wall. GALib: A C++ Library of Genetic Algorithm Components. Version 2.4, Documentation Revision B. Technical Report, Mechanical Engineering Department, Massachusetts Institute of Technology, August 1996. Online available at <http://lancet.mit.edu/ga/dist/galibdoc.pdf> [accessed 2007-08-22].
- [1690] Maarten Keijzer, J. J. Merelo, G. Romero, and Marc Schoenauer. Evolving Objects: A General Purpose Evolutionary Computation Library. In *Proceedings of 5th International Conference on Artificial Evolution, Evolution Artificielle, EA 2001*, pages 829–888, 2001. In proceedings [513]. Online available at <http://www.lri.fr/~marc/EO/EO-EA01.ps.gz> and <http://citeseer.ist.psu.edu/keijzer01evolving.html> [accessed 2007-08-24]. See also <http://eodev.sourceforge.net/> [accessed 2007-08-24].
- [1691] Evelyne Lutton, Pierre Collet, and Jean Louchet. EASEA Comparisons on Test Functions: GALib versus EO. In *Proceedings of the Fifth Conference on Artificial Evolution, Evolution Artificielle (EA-2001)*, pages 219–230, 2001. In proceedings [513]. Online available at <http://fractales.inria.fr/evo-lab/EASEAComparisonFinal.ps.gz> [accessed 2007-08-24].
- [1692] James Gosling and Henry McGilton. The Java Language Environment – A White Paper. Technical Report, Sun Microsystems, Inc., May 1996. Online available at <http://java.sun.com/docs/white/langenv/> [accessed 2007-07-03].
- [1693] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *Java(TM) Language Specification*. The Java Series. Prentice Hall PTR, 3rd edition, June 2005. ISBN 978-0321246783. Online available at <http://java.sun.com/docs/books/jls/> [accessed 2007-09-14].
- [1694] Jon Byous. Java Technology: The Early Years. Technical Report, Sun Microsystems, Inc., ca. 1998. Online available at <http://java.sun.com/features/1998/05/birthday.html> [accessed 2007-07-03].
- [1695] Guido Krüger. *Handbuch der Java-Programmierung*. Addison-Wesley, 4. aktualisierte edition, 2006. ISBN 3-8273-2361-4 and 3-8273-2447-5. Online available at <http://www.javabuch.de/> [accessed 2007-07-03].
- [1696] Bruce Eckel. *Thinking in Java*. Prentice Hall PTR, 4th edition, February 2006. ISBN 978-0131872486. 3rd edition online available at <http://www.mindview.net/Books/TIJ/> [accessed 2007-07-03].
- [1697] David Flanagan. *Java In A Nutshell*. O'Reilly Media, Inc., 5th edition, March 2005. ISBN 978-0596007737.
- [1698] David Flanagan. *Java Examples in a Nutshell*. O'Reilly Media, Inc., 3rd, illustrated edition, January 2004. ISBN 978-0596006204.
- [1699] Pat Niemeyer and Jonathan Knudsen. *Learning Java*. O'Reilly Media, Inc, third edition, June 2005. ISBN 978-0-596-00873-4.
- [1700] Christian Ullenboom. *Java ist auch eine Insel – Programmieren mit der Java Standard Edition Version 6*. Galileo Computing, Galileo Press, 6. aktualisierte und erweiterte edition, 2007. ISBN 3-89842-838-9. Online available at <http://www.galileocomputing.de/openbook/javainssel6/> [accessed 2007-07-03].
- [1701] Yiping Ding, Ethan D. Bolker, and Arjun Kumar. Performance Implications of Hyper-Threading. In *Proceedings of 29th International Computer Measurement Group Conference*, pages 21–29. Computer Measurement Group, December 7–12, 2003, Dallas, Texas, USA.
- [1702] Marcel Dekker. *Introduction to Set Theory*. CRC, revised, and expanded, third edition, June 22, 1999. ISBN 0-8247-7915-0.
- [1703] Paul R. Halmos. *Naive Set Theory*. Undergraduate Texts in Mathematics. Springer-Press New York Inc., first edition, June 2001. ISBN 0-3879-0092-6. New Edition January 1998, ISBN: 978-0387900926.
- [1704] Robert R. Stoll. *Set Theory and Logic*. Dover Publications, reprint edition, October 1, 1979. ISBN 0-4866-3829-4, 978-0486638294.

- [1705] Donald Ervin Knuth. *Sorting and Searching*, volume 3 of *The Art of Computer Programming (TAOCP)*. Addison-Wesley, Reading, Massachusetts, second edition, 1998. ISBN 0-201-89685-0.
- [1706] Robert Sedgewick. *Algorithms in Java, Parts 1-4*. Addison Wesley, third edition, September 2002. ISBN 0-2013-6120-5.
- [1707] Steven R. Finch. Transitive Relations, Topologies and Partial Orders, June 5, 2003. Online available at <http://citeseer.ist.psu.edu/finch03transitive.html> and <http://algo.inria.fr/csolve/posets.pdf> [accessed 2007-07-28]. See also [1960].
- [1708] Bas C. van Fraassen. *Laws and Symmetry*. Clarendon Paperbacks. Oxford University Press, USA, January 1990. ISBN 978-0198248606.
- [1709] David A. R. Wallace. *Groups, Rings and Fields*. Springer Undergraduate Mathematics Series. Springer, February 2004. ISBN 978-3540761778. Online available at <http://books.google.de/books?id=EmO9ejuMHNuc> [accessed 2007-07-28].
- [1710] Alfréd Rényi. *Probability Theory*. Dover Publications, May 2007. ISBN 978-0486458670.
- [1711] Gregory F. Lawler. *Introduction to Stochastic Processes*. Chapman & Hall/CRC, second edition, May 2006. ISBN 978-1584886518.
- [1712] Edwin Thompson Jaynes. *Probability Theory: The Logic of Science*. Washington University, preprint edition, 1996. See also printed version [1713]. G. Larry Bretthorsted. Online available at <http://bayes.wustl.edu/etj/prob/book.pdf> [accessed 2007-08-08].
- [1713] Edwin Thompson Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, June 2003. ISBN 978-0521592710. See also [1712]. Larry Bretthorsted. Online available at <http://bayes.wustl.edu/etj/prob/book.pdf> [accessed 2007-08-08].
- [1714] Olav Kallenberg. *Foundations of modern probability*. Probability and Its Applications. Springer, New York, second edition, January 8, 2002. ISBN 978-0387953137, 0-3872-5115-4. OCLC: 60740995.
- [1715] Olav Kallenberg. *Probabilistic symmetries and invariance principles*. Probability and its Applications. Springer, New York, July 27, 2005. ISBN 978-0-387-25115-8, 0-3879-4957-7, 0-3879-5313-2. OCLC: 46937587.
- [1716] Pierre-Simon Marquis de Laplace. *Théorie Analytique des Probabilités (Analytical Theory of Probability)*. Courcier, Imprimeur-Libraire pour les Mathématiques, quai des Augustins, no. 57, 1812. Première Partie. Online available at <http://books.google.de/books?id=nQwAAAAAMAAJ> [accessed 2007-08-27].
- [1717] Richard von Mises. *Probability, Statistics, and Truth*. Dover Publications, second english, reproduced 1981 dover edition, 1957. ISBN 0486242145, 978-0486242149.
- [1718] Richard von Mises and Hilda Geiringer. *Mathematical Theory of Probability and Statistics*. Academic Press Inc., New York, February 1965. ISBN 0127273565, 978-0127273563.
- [1719] Andrei Nikolajevich Kolmogorov. *Foundations of the Theory of Probability*. Chelsea Publishing Company New York, second edition, 1956, June 1960. ISBN 978-0828400237. monograph “Grundbegriffe der Wahrscheinlichkeitsrechnung” 1933, book 1950, Online available at <http://www.mathematik.com/Kolmogorov/> [accessed 2007-09-15].
- [1720] Paul T. von Hippel. Mean, Median, and Skew: Correcting a Textbook Rule. *Journal of Statistics Education*, 13(2), 2005. Online available at <http://www.amstat.org/publications/jse/v13n2/vonhippel.html> [accessed 2007-09-15].
- [1721] Henry Bottomley. Relationship between the mean, median, mode, and standard deviation in a unimodal distribution, September 2006. Online available at <http://www.btinternet.com/~se16/hgb/median.htm> [accessed 2007-09-15].
- [1722] Claude Elwood Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:379–423/623–656, July/October 1948. Online available at

- <http://plan9.bell-labs.com/cm/ms/what/shannonday/paper.html> [accessed 2007-09-15]. Also published in D. Slepian, editor, *Key Papers in the Development of Information Theory*, New York: IEEE Press, 1974; N. J. A. Sloane and A. D. Wyner, editors, *Claude Elwood Shannon: Collected Papers*, New York: IEEE Press, 1993; W. Weaver and Claude Elwood Shannon, *The Mathematical Theory of Communication*, Urbana, Illinois: University of Illinois Press, 1949, republished in paperback 1963.
- [1723] A. C. G. Verdugo Lazo and P. N. Rathie. On the entropy of continuous probability distributions. *IEEE Transactions on Information Theory*, 24(1):120–122, 1978.
 - [1724] J. H. Ahrens and U. Dieter. Computer methods for sampling from gamma, beta, poisson and binomial distributions. *Computing*, 12(3):223–246, September 10, 1974. ISSN 0010-485X (Print) 1436-5057 (Online). doi:10.1007/BF02293108. Online available at <http://www.springerlink.com/content/712tn58716485674/fulltext.pdf> [accessed 2007-09-15].
 - [1725] Donald L. Snyder and Michael I. Miller. *Random Point Processes in Time and Space*. Springer, second edition, June 19, 1991. ISBN 978-0387975771.
 - [1726] Marvin Zelen and Norman C. Severo. Probability Functions. In Milton Abramowitz and Irene A. Stegun, editors, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, chapter 26. Dover Publications / National Bureau of Standards, first. (new ed june 1, 1965) edition, 1964. ISBN 978-0486612720.
 - [1727] Ming Tan, Hong-Bin Fang, Guo-Liang Tian, and Gang Wei. Testing multivariate normality in incomplete data of small sample size. *Journal of Multivariate Analysis*, 93(1):164–179, 2005. ISSN 0047-259X. Online available at <http://dx.doi.org/10.1016/j.jmva.2004.02.014> [accessed 2007-09-15].
 - [1728] N. J. H. Small. Miscellanea: Marginal Skewness and Kurtosis in Testing Multivariate Normality. *Applied Statistics*, 29(1):85–87, 1980. ISSN 0035-9254.
 - [1729] J. P. Royston. Some Techniques for Assessing Multivariate Normality Based on the Shapiro-Wilk W. *Applied Statistics*, 32(2):121–133, 1983. ISSN 0035-9254.
 - [1730] Luc Devroye from the School of Computer Science, McGill University. *Non-Uniform Random Variate Generation*. Springer-Verlag, New York, 1986. ISBN 0-387-96305-7, 3-540-96305-7. Online available at <http://cg.scs.carleton.ca/~luc/rnbookindex.html> [accessed 2007-07-05].
 - [1731] William Sealy Gosset. The probable error of a mean. *Biometrika*, 6(1):1–25, March 1908. Because the author was not allowed to publish this article, he used the pseudonym *Student*. Online available at <http://www.york.ac.uk/depts/maths/histstat/student.pdf> [accessed 2007-09-30]. Reprinted in [1732].
 - [1732] William Sealy Gosset. The probable error of a mean. In E. S. Pearson and John Wishart, editors, “*Student’s*” *Collected Papers*, pages 11–34. Cambridge University Press for the Biometrika Trustees, 1942. With a Foreword by Launce McMullen. Reprint of [1731].
 - [1733] Sir Ronald Aylmer Fisher. Applications of “Student’s” distribution. *Metron*, 5:90–104, 1925. Online available at <http://digital.library.adelaide.edu.au/coll/special/fisher/43.pdf> [accessed 2007-09-30].
 - [1734] John A. Rice. *Mathematical Statistics and Data Analysis*. Statistics. Duxbury Press, third edition, April 2006. ISBN 978-0534399429, 978-0534209346 (2nd ed.).
 - [1735] Steven M. Kay. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*, volume 1. Prentice Hall PTR, us es edition, March 26, 1993. ISBN 978-0133457117.
 - [1736] H. Vincent Poor. *An Introduction to Signal Detection and Estimation*. Springer, second edition, June 2005. ISBN 978-0387941738.
 - [1737] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I*. Wiley-Interscience, reprint edition, February 2007. ISBN 978-0471095170.
 - [1738] Dan Simon. *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*. Wiley-Interscience, June 2006. ISBN 978-0471708582.

- [1739] John Aldrich. R.A. Fisher and the making of maximum likelihood 1912-1922. *Statistical Science*, 3(12):162–176, January 1995. Online available at <http://projecteuclid.org/euclid.ss/1030037906> [accessed 2007-09-15].
- [1740] R. L. Plackett. Some theorems in least squares. *Biometrika*, 37(12):149–157, 1950. Online available at <http://biomet.oxfordjournals.org/cgi/reprint/37/1-2/149> [accessed 2007-09-15].
- [1741] Sir Ronald Aylmer Fisher. *Statistical methods and scientific inference*. New York, Hafner Press/Macmillan Pub Co, revised (june 1973) edition, 1956. ISBN 0028447409.
- [1742] George Casella and Roger L. Berger. *Statistical Inference*. Duxbury Thomson Learning, Pacific Grove, CA, USA, second edition, 2002. ISBN 0-534-24312-6.
- [1743] M.C. Jones Bernard. W. Silverman. Fix, E. and Hodges, J. L. (1951): An important contribution to nonparametric discriminant analysis and density estimation: Commentary on Fix and Hodges (1951). *International Statistical Review / Revue Internationale de Statistique*, 57(3):233–247, April 1989.
- [1744] David W. Scott. *Multivariate Density Estimation: Theory, Practice, and Visualization*. Wiley Series in Probability and Statistics. Wiley-Interscience, John Wiley & Sons, August 1992. ISBN 0-4715-4770-0.
- [1745] Bernard. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, Bristol, England, April 1986. ISBN 0-4122-4620-1.
- [1746] Kalyanmoy Deb, Samir Agrawal, Amrit Pratab, and T. Meyarivan. A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II. In *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pages 849–858, 2000. In proceedings [587]. Online available at <http://citeseer.ist.psu.edu/deb00fast.html> [accessed 2007-07-28].
- [1747] Emanuel Parzen from the Stanford University California, Applied Mathematics and Statistics Lab. On estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33:1065–1076, August 1962. Online available at <http://citeseer.ist.psu.edu/parzen62estimation.html> [accessed 2007-08-11].
- [1748] Miloš Drutarovský, Viktor Fischer, Martin Šimka, and Frédéric Celle. A Simple PLL-based True Random Number Generator for Embedded Digital Systems. *Computing and Informatics*, 23(5):501–515, 2004.
- [1749] Martin Šimka, Miloš Drutarovský, and Viktor Fischer. Embedded True Random Number Generator in Actel FPGAs. In *Workshop on Cryptographic Advances in Secure Hardware – CRASH 2005*, pages 6–7, September 2005. Leuven, Belgium. Online available at <http://www.best.tuke.sk/simka/pub.html> [accessed 2007-09-15].
- [1750] Viktor Fischer, Miloš Drutarovský, Martin Šimka, and Nathalie Bochard. High Performance True Random Number Generator in Altera Stratix FPLDs. In Jürgen Becker, Marco Platzner, and Serge Vernalde, editors, *Field-Programmable Logic and Applications – FPL 2004*, volume 3203 of *Lecture Notes in Computer Science (LNCS)*, pages 555–564, August 2004. Springer-Verlag, Lueven, Belgium. Online available at <http://www.best.tuke.sk/simka/pub.html> [accessed 2007-09-15].
- [1751] Berk Sunar, William J. Martin, and Douglas R. Stinson. A Provably Secure True Random Number Generator with Built-In Tolerance to Active Attacks. *IEEE Transactions on Computers*, 56(1):109–119, 2007. ISSN 0018-9340. doi:10.1109/TC.2007.4. Online available at <http://www.cacr.math.uwaterloo.ca/~dstinson/papers/rng-IEEE.pdf> [accessed 2007-09-15].
- [1752] Martin Šimka. Testing True Random Number Generators Used in Cryptography. In *Proceedings of the IV. PhD students conference*, pages 95–96. Technical University of Košice, Slovakia, May 2004. Online available at <http://www.best.tuke.sk/simka/pub.html> [accessed 2007-09-15].
- [1753] Ueli Maurer. Fast Generation of Prime Numbers and Secure Public-Key Cryptographic Parameters. *Journal of Cryptology*, 8(3):123–155, September 1995. ISSN 0933-2790 (Print) 1432-1378 (Online). doi:10.1007/BF00202269. Online avail-

- able at <http://www.springerlink.com/content/u3710818146qq153/fulltext.pdf> and <http://citeseer.ist.psu.edu/186041.html> [accessed 2007-09-15].
- [1754] Ilpo Vattulainen, T. Ala-Nissila, and K. Kankaala. Physical tests for random numbers in simulations. *Physical Review Letters*, 73(7):2513–2516, 1994. Online available at <http://citeseer.ist.psu.edu/vattulainen94physical.html> [accessed 2007-07-28].
 - [1755] Walter Selke, L.N. Shchur, and A.L. Talapov. Cluster-flipping Monte Carlo algorithm and correlations in “good” random number generators. *JETP Letters*, 58(8):684–686, 1993.
 - [1756] Alan M. Ferrenberg, D. P. Landau, and Y. Joanna Wong. Monte Carlo simulations: Hidden errors from “good” random number generators. *Physical Review Letters*, 69(23):3382–3384, December 1992. Online available at http://prola.aps.org/pdf/PRL/v69/i23/p3382_1 [accessed 2007-08-18].
 - [1757] Hui-Chin Tang. Combined random number generator via the generalized Chinese remainder theorem. *Journal of Computational and Applied Mathematics*, 142(2):377–388, May 15, 2002. ISSN 0377-0427. doi:10.1016/S0377-0427(01)00424-1. Online available at [http://dx.doi.org/10.1016/S0377-0427\(01\)00424-1](http://dx.doi.org/10.1016/S0377-0427(01)00424-1) [accessed 2007-09-16].
 - [1758] Abhishek Parakh from the Louisiana State University. A d-Sequence based Recursive Random Number Generator, 2006. Online available at <http://arxiv.org/abs/cs.CR/0603029> [accessed 2007-09-16].
 - [1759] Lenore Blum, Manuel Blum, and Michael Shub. A simple unpredictable pseudo-random number generator. *SIAM Journal on Computing*, 15:364–383, May 1986. ISSN 0097-5397.
 - [1760] E. F. Carter. The Generation and Application of Random Numbers. *Forth Dimensions*, XVI(1 and 2), 1994, Oakland California.
 - [1761] Michael Henneke from the Rechenzentrum, Universität Karlsruhe, D-76128 Karlsruhe, Germany. RANEXP: experimental random number generator package. *Computer Physics Communications*, 79(2):261–267, April 1994. doi:10.1016/0010-4655(94)90072-8. Online available at [http://dx.doi.org/10.1016/0010-4655\(94\)90072-8](http://dx.doi.org/10.1016/0010-4655(94)90072-8) [accessed 2007-08-19].
 - [1762] Claude Overstreet, Jr. and Richard E. Nance. A random number generator for small word-length computers. In *ACM’73: Proceedings of the annual conference*, pages 219–223, 1973, Atlanta, Georgia, United States. ACM Press/CSC-ER, New York, NY, USA. doi:10.1145/800192.805707. Online available at <http://doi.acm.org/10.1145/800192.805707> [accessed 2007-09-15].
 - [1763] P. D. Coddington and A. J. Newell. JAPARA – A Java Parallel Random Number Generator Library for High-Performance Computing. In *Proceedings of the 6th International Workshop on Java for Parallel and Distributed Computing IPDPS 2004*, page 156a. IEEE Computer Society, Los Alamitos, CA, USA, April 26, 2004, Santa Fe, New Mexico. ISBN 076-9521-320. doi:10.1109/IPDPS.2004.1303143.
 - [1764] John R. Koza from the Computer Science Department, Stanford University, Stanford, CA 94305 USA. Evolving a computer program to generate random numbers using the genetic programming paradigm. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 37–44, 1991. In proceedings [749]. Online available at <http://citeseer.ist.psu.edu/112986.html> and <http://www.genetic-programming.com/jkpdf/icga1991.pdf> [accessed 2008-05-29].
 - [1765] Derrick Henry Lehmer. Mathematical methods in large-scale computing units. In *Proceedings of the 2nd Symposium on Large-Scale Digital Calculating Machinery, 1949*, pages 141–146. Harvard University Press, 1951, Cambridge, MA, USA.
 - [1766] Derrick Henry Lehmer. Mathematical methods in large-scale computing units. *Math. Rev.*, 13(1):495, 1952.
 - [1767] Karl Entacher. Bad subsequences of well-known linear congruential pseudorandom number generators. *ACM Transactions on Modeling and Computer Simulation*

- (*TOMACS*), 8(1):61–70, 1998. ISSN 1049-3301. doi:10.1145/272991.273009. Online available at <http://doi.acm.org/10.1145/272991.273009> [accessed 2007-09-15].
- [1768] Donald Ervin Knuth. *The Art of Computer Programming: Seminumerical Algorithms*, volume 2. Addison-Wesley, third edition, 1997. ISBN 020-1896-842, 978-0201038026.
- [1769] George Edward Pelham Box and Mervin Edgar Muller. A note on the generation of random normal deviates. *Annals Math. Stat.*, 29(5):610–611, 1958. Online available at <http://projecteuclid.org/euclid.aoms/1177706645> [accessed 2007-09-15].
- [1770] *Classification, Clustering, and Data Mining Applications: Proceedings of the Meeting of the International Federation of Classification Societies (IFCS), ... Data Analysis, and Knowledge Organization*, 2004. Springer-Verlag New York, Inc., Secaucus, NJ, USA. ISBN 3540220143.
- [1771] Margaret H. Dunham. *Data Mining: Introductory and Advanced Topics*. Prentice Hall PTR, Upper Saddle River, NJ, USA, August 2002. ISBN 0130888923.
- [1772] Boris Mirkin. *Clustering for Data Mining: A Data Recovery Approach*. Computer Science and Data Analysis. Chapman & Hall/CRC, April 2005. ISBN 978-1584885344.
- [1773] Pavel Berkhin. Survey Of Clustering Data Mining Techniques. Technical Report, Accrue Software, San Jose, CA, 2002. Online available at http://www.ee.ucr.edu/~barth/EE242/clustering_survey.pdf and <http://citeseer.ist.psu.edu/berkhin02survey.html> [accessed 2007-08-27].
- [1774] Michael W. Berry, editor. *Survey of Text Mining: Clustering, Classification, and Retrieval*. Springer, September 2003. ISBN 978-0387955636.
- [1775] Mieczysław A. Kłopotek, Sławomir T. Wieruchoń, and Krzysztof Trojanowski, editors. *Intelligent Information Processing and Web Mining: Proceedings of the International IIS: IIPWMA05*, Advances in Soft Computing, June 2005, Gdansk, Poland. Springer, Berlin, Heidelberg, New York. ISBN 978-3-540-25055-2.
- [1776] Mingfang Wu, Michael Fuller, and Ross Wilkinson. Using clustering and classification approaches in interactive retrieval. *Inf. Process. Manage.*, 37(3):459–484, 2001. Online available at <http://citeseer.ist.psu.edu/wu01using.html> and <http://de.scientificcommons.org/313591> [accessed 2007-08-11].
- [1777] Dmitri Roussinov and Hsinchun Chen. Information navigation on the web by clustering and summarizing query results. *Information Processing & Management*, 37(6):789–816, 2001. Online available at [http://dx.doi.org/10.1016/S0306-4573\(00\)00062-5](http://dx.doi.org/10.1016/S0306-4573(00)00062-5) [accessed 2007-08-11].
- [1778] James C. Bezdek, James Keller, Raghu Krishnapuram, and Nikhil R. Pal. *Fuzzy Models and Algorithms for Pattern Recognition and Image Processing*. The Handbooks of Fuzzy Sets. Springer, March 2005. ISBN 978-0387245157.
- [1779] C. S. Warnekar and G. Krishna. A heuristic clustering algorithm using union of overlapping pattern-cells. *Pattern Recognition*, 11(2):85–93, 1979. Online available at [http://dx.doi.org/10.1016/0031-3203\(79\)90054-2](http://dx.doi.org/10.1016/0031-3203(79)90054-2) [accessed 2007-08-11]. (Link not viable on 2007-08-27).
- [1780] Sanjeev Jagannatha Koppal and Srinivasa G Narasimhan. Clustering Appearance for Scene Analysis. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1323–1330, June 2006. Online available at http://www.ri.cmu.edu/pubs/pub_5376.html [accessed 2007-08-11].
- [1781] Bruce J. Schachter, Larry S. Davis, and Azriel Rosenfeld. Some experiments in image segmentation by clustering of local feature values. *Pattern Recognition*, 11(1):19–28, 1979.
- [1782] Jung Eun Shim and Won Suk Lee. A landmark extraction method for protein 2DE gel images based on multi-dimensional clustering. *Artificial Intelligence in Medicine*, 35(1-2):157–170, 2005. Online available at <http://linkinghub.elsevier.com/retrieve/pii/S0933366570500076X> and <http://dx.doi.org/10.1016/j.artmed.2005.07.002> [accessed 2007-08-11].

- [1783] F. A. da Veiga. Structure discovery in medical databases: a conceptual clustering approach. *Artificial Intelligence in Medicine*, 8(5):473–491, 1996. Online available at [http://dx.doi.org/10.1016/S0933-3657\(96\)00353-3](http://dx.doi.org/10.1016/S0933-3657(96)00353-3) [accessed 2007-08-11].
- [1784] Song Zhang and David H. Laidlaw. DTI Fiber Clustering and Cross-subject Cluster Analysis. In *Proceedings International Society for Magnetic Resonance in Medicine (ISMRM)*, May 2005. Miami, FL. Online available at <http://www.cs.brown.edu/research/vis/docs/pdf/Zhang-2005-DFC.pdf> [accessed 2007-08-11].
- [1785] Anil K. Jain, M. N. Murty, and P. J. Flynn. Data Clustering: A Review. *ACM Computing Surveys*, 31(3):264–323, September 1999. Online available at <http://citeseer.ist.psu.edu/jain99data.html> and <http://www.cs.rutgers.edu/~mlittman/courses/lightai03/jain99data.pdf> [accessed 2007-08-11].
- [1786] Uzay Kaymak and Magne Setnes. Extended Fuzzy Clustering Algorithms. Research Paper ERS; ERS-2000-51-LIS, Erasmus Research Institute of Management (ERIM), RSM Erasmus University, November 2000. Online available at <https://ep.eur.nl/handle/1765/57> and <http://ideas.repec.org/p/dgr/eureri/200050.html> [accessed 2007-08-11].
- [1787] R. Krishnapuram, A. Joshi, and L. Yi O. Nasraoui. Low-Complexity Fuzzy Relational Clustering Algorithms for Web Mining. *IEEE-FS*, 9:595–607, August 2001. Online available at <http://citeseer.ist.psu.edu/krishnapuram01lowcomplexity.html> and <http://de.scientificcommons.org/583343> [accessed 2007-08-11].
- [1788] Steffen Bleul. Ähnlichkeitsmaße und Clustering (Text-based Similarity Measures and Clustering), 2002. Online available at <http://www.vs.uni-kassel.de/~bleul/> [accessed 2007-08-11].
- [1789] Hervé Abdi. Centroid, Center of Gravity, Center of Mass, Barycenter. In Neil J. Salkind, editor, *Encyclopedia of Measurement and Statistics*, volume 1, page 128 ff. Thousand Oaks (California): Sage Publications, Inc, October 2006. ISBN 978-1412916110. Online available at <http://www.utdallas.edu/~herve/Abdi-Centroid2007-pretty.pdf> [accessed 2007-08-11].
- [1790] Anil K. Jain and Richard C. Dubes. *Algorithms for Clustering Data*. Prentice-Hall Advanced Reference Series. Prentice-Hall, Upper Saddle River, NJ, USA, 1988. ISBN 0-13-022278-X.
- [1791] J. B. MacQueen. Some Methods for Classification and Analysis of Multivariate Observations. In *Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297. Berkeley, University of California Press, 1967.
- [1792] Sergei Vassilvitskii. How Slow is the k-Means Method? *Discrete and Computational Geometry*, (Special Issue: Proceedings of 22nd Annual ACM Symposium on Computational Geometry), June 2006, Sedona, Arizona. Online available at <http://www.stanford.edu/~sergeiv/papers/kMeans-socg.pdf> and <http://portal.acm.org/citation.cfm?id=1137880> [accessed 2007-08-11].
- [1793] Eric W. Weisstein. K-Means Clustering Algorithm, 1999–2006. From MathWorld—A Wolfram Web Resource. Online available at <http://mathworld.wolfram.com/K-MeansClusteringAlgorithm.html> [accessed 2007-08-11].
- [1794] Eckart Zitzler and Lothar Thiele. An Evolutionary Algorithm for Multiobjective Optimization: The Strength Pareto Approach. Technical Report 43, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology Zürich (ETH), Gloriastrasse 35, CH-8092 Zurich, Switzerland, May 1998. Online available at <http://www.tik.ee.ethz.ch/sop/publicationListFiles/zit1998a.pdf> and <http://citeseer.ist.psu.edu/225338.html> [accessed 2007-07-29].
- [1795] S. Asharafa and M. Narasimha Murtyb. An adaptive rough fuzzy single pass algorithm for clustering large data sets. *Pattern recognition (Pattern recogn.)*, 36(12):3015–3018, 2003. ISSN 0031-3203.

- [1796] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison Wesley, international ed edition, January 1974. ISBN 978-0201000290.
- [1797] Mikhail J. Atallah and Susan Fox, editors. *Algorithms and Theory of Computation Handbook*. CRC Press, Inc., Boca Raton, FL, USA, first edition, November 1998. ISBN 978-0849326493. Produced By Suzanne Lassandro.
- [1798] Sara Baase and Allen Van Gelder. *Computer Algorithms: Introduction to Design and Analysis*. Addison Wesley, third edition, November 1999. ISBN 978-0201612448.
- [1799] Dexter C. Kozen. *The Design and Analysis of Algorithms (Texts & Monographs in Computer Science)*. Springer New York, January 1992. ISBN 978-0387976877.
- [1800] Arthur W. Burks. From ENIAC to the Stored-Program Computer: Two Revolutions in Computers. In N. Metropolis, J. Howlett, and Gian-Carlo Rota, editors, *A History of Computing in the Twentieth Century: A Collection of Essays with Introductory Essay and Indexes*, pages 311–344. Academic Press, New York, November 1980. ISBN 978-0124916500. Papers from the Los Alamos International Research Conference on the History of Computing, 1976.
- [1801] David A. Patterson and John L. Hennessy. *Computer Organization and Design: The Hardware/Software Interface*. The Morgan Kaufmann Series in Computer Architecture and Design. Morgan Kaufmann, third edition, August 2004. ISBN 978-1558606043.
- [1802] David Salomon. *Assemblers and loaders*. Ellis Horwood Series in Computers and Their Applications. Ellis Horwood, Upper Saddle River, NJ, USA, February 1993. ISBN 978-0130525642, 0-13-052564-2.
- [1803] Masud Ahmad Malik from the Department of Computer Science & Information Systems, Philadelphia University, P.O. Box 1101, Sweileh. Amman, Jordan. Evolution of the high level programming languages: a critical perspective. *ACM SIGPLAN Notices*, 33(12):72–80, December 1998. ISSN 0362-1340. Online available at <http://doi.acm.org/10.1145/307824.307882> [accessed 2007-09-14].
- [1804] Donald Ervin Knuth. *Fundamental Algorithms*, volume 1 of *The Art of Computer Programming (TAOCP)*. Addison-Wesley, Reading, Massachusetts, third edition, 1997. ISBN 0-201-89683-4.
- [1805] Ingo Wegener and Randall Pruim (Translator). *Complexity Theory: Exploring the Limits of Efficient Algorithms*. Springer-Verlag, Berlin Heidelberg, 2005. ISBN 978-3540210450. Translated from the German “Komplexitätstheorie – Grenzen der Effizienz von Algorithmen”, (Springer-Verlag 2003, ISBN: 3-540-00161-1).
- [1806] Herbert S. Wilf. *Algorithms and Complexity*. AK Peters, Ltd., second edition, December 2002. ISBN 978-1568811789.
- [1807] Donald E. Knuth. Big Omicron and big Omega and big Theta. *SIGACT News*, 8(2):18–24, 1976. ISSN 0163-5700. Online available at <http://portal.acm.org/citation.cfm?id=1008328.1008329> [accessed 2007-08-11].
- [1808] Sir Charles Antony Richard (Tony) Hoare. Quicksort. *Computer Journal*, 5(1):10–15, 1962.
- [1809] Paul Gustav Heinrich Bachmann. *Die Analytische Zahlentheorie / Dargestellt von Paul Bachmann*, volume Zweiter Theil of *Zahlentheorie : Versuch einer Gesamtdarstellung dieser Wissenschaft in ihren Haupttheilen*. B. G. Teubner, Leipzig, Germany, 1894. Online available at <http://gallica.bnf.fr/ark:/12148/bpt6k994750> [accessed 2008-05-20]. Republished by Scholarly Publishing Office, University of Michigan Library, ISBN 141818540X, 978-1418185404, 2006.
- [1810] Edmund Landau. *Handbuch der Lehre von der Verteilung der Primzahlen*. B. G. Teubner, Leipzig, Germany, 1909. Reprinted by Chelsea, New York, 1953.
- [1811] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge International Series on Parallel Computation. Cambridge University Press, August 1995. ISBN 978-0521474658.

- [1812] Juraj Hromkovic. *Algorithmics for Hard Computing Problems*. Springer, first edition, June 2001. ISBN 978-3540668602.
- [1813] Michael Mitzenmacher and Eli Upfal. *Probability and Computing : Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, January 2005. ISBN 0521835402.
- [1814] J. Hromkovic and I. Zámecniková. *Design and Analysis of Randomized Algorithms: Introduction to Design Paradigms*. Texts in Theoretical Computer Science, an EATCS series. Springer, first edition, July 2005. ISBN 978-3540239499.
- [1815] Rajeev Motwani and Prabhakar Raghavan. Randomized algorithms. *ACM Comput. Surv.*, 28(1):33–37, 1996. ISSN 0360-0300.
- [1816] Harvey Gould and Jan Tobochnik. *An Introduction to Computer Simulation Methods: Part 2, Applications to Physical Systems*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995. ISBN 0201506041, 978-0201165036. Edited by Julia Berrisford.
- [1817] P. K. Mackeown. *Stochastic Simulation in Physics*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2001. ISBN 9813083263.
- [1818] Christian P. Robert and George Casella. *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005. ISBN 0387212396.
- [1819] Jun S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, October 2002. ISBN 0387952306.
- [1820] Gero Mühl from the Kommunikations- und Betriebssysteme, Fakultät für Elektrotechnik und Informatik, Technische Universität Berlin. Lecture Notes “Verteilte Systeme”, Lesson “Einführung und Motivation, 2006. Online available at <http://kbs.cs.tu-berlin.de/teaching/sose2006/vs/folien/02.pdf> [accessed 2008-02-09].
- [1821] Friedemann Mattern. *Verteilte Basisalgorithmen*. Springer-Verlag GmbH, 1989. ISBN 978-3540518358.
- [1822] David G. Andersen, Hari Balakrishnan, Frans Kaashoek, and Robert Morris. Resilient Overlay Networks. In *18th ACM SOSP*, October 2001. Banff, Canada. Online available at <http://nms.csail.mit.edu/ron/> and <http://nms.lcs.mit.edu/papers/ron-sosp2001.pdf> [accessed 2007-08-13].
- [1823] Wanlei Zhou and Andrzej Goscinski. An Analysis of the Web-Based Client-Server Computing Models. In *Proceedings of the Asia-Pacific Web Conference (APWeb'98)*, pages 343–348, September 1998, Beijing. Online available at <http://citeseer.ist.psu.edu/296101.html> [accessed 2007-08-13].
- [1824] G. Degli Antoni, D. Cabianca, M. Vaccari, M. Benini, and F. Casablanca. Linearity of Client/Server Systems. *Bulletin of the European Association for Theoretical Computer Science*, 57:201–214, April 1995. Online available at <http://citeseer.ist.psu.edu/antoni95linearity.html> and <http://en.scientificcommons.org/130995> [accessed 2007-08-13].
- [1825] George F. Coulouris, Jean Dollimore, and Tim Kindberg. *Distributed Systems: Concepts and Design*. Addison Wesley, fourth rev. edition, June 2005. ISBN 978-0321263544.
- [1826] Andrew S. Tanenbaum and Maarten van Steen. *Distributed Systems. Principles and Paradigms*. Prentice Hall International, international ed edition, March 2003. ISBN 978-0131217867. Some information available at <http://www.cs.vu.nl/~ast/books/ds1/> [accessed 2007-08-13].
- [1827] Ralf Steinmetz and Klaus Wehrle. Peer-to-Peer-Networking & -Computing – Aktuelles Schlagwort. *Informatik Spektrum*, 27(1):51–54, 2004. Online available at <http://www.springerlink.com/content/up3vdx3cnu1a4wb3/fulltext.pdf> [accessed 2007-08-13].
- [1828] Luc Onana Alima, Sameh El-Ansary, Per Brand, and Seif Haridi. $\mathcal{DKS}(\mathcal{N}, k, f)$: A Family of Low Communication, Scalable and Fault-Tolerant Infrastructures for

- P2P Applications. In *3rd IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2003)*, pages 344–350. IEEE Computer Society, 2003. ISBN 0-7695-1919-9. Online available at <http://dks.sics.se/pub/ccgrid-dks.pdf> [accessed 2007-08-13].
- [1829] Stephanos Androutsellis-Theotokis and Diomidis Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Comput. Surv.*, 36(4):335–371, 2004. ISSN 0360-0300. Online available at <http://www.spinellis.gr/pubs/jrnl/2004-ACMCS-p2p/html/AS04.pdf> and <http://portal.acm.org/citation.cfm?id=1041681> [accessed 2007-08-13].
- [1830] Stefan Saroiu, P. Krishna Gummadi, and Steven Gribble. A Measurement Study of Peer-to-Peer File Sharing Systems. In *SPIE Multimedia Computing and Networking (MMCN2002)*, January 2002, San Jose, CA, USA. Online available at <http://www.cs.washington.edu/homes/gribble/papers/mmcn.pdf> and <http://citeseer.ist.psu.edu/saroiu02measurement.html> [accessed 2007-08-13].
- [1831] Ivan Stojmenović, editor. *Handbook of Sensor Networks: Algorithms and Architectures*. Wiley-Interscience, October 5, 2005. ISBN 978-0-471-68472-5.
- [1832] David Culler, Deborah Estrin, and Mani Srivastava. Guest Editors' Introduction: Overview of Sensor Networks. *Computer*, 37(8):41–49, August 2004. ISSN 0018-9162. Online available at <http://www.archrock.com/downloads/resources/IEEE-overview-2004.pdf> [accessed 2007-09-15].
- [1833] S. Sitharama Iyengar and Richard R. Brooks, editors. *Distributed Sensor Networks*. Chapman & Hall/CRC, December 29, 2004. ISBN 978-1584883838.
- [1834] Cauligi S. Raghavendra, Krishna M. Sivalingam, and Taieb Znati, editors. *Wireless Sensor Networks (ERCOTAC)*. Springer Netherlands, second edition, June 1, 2004. ISBN 978-1402078835.
- [1835] Feng Zhao and Leonidas Guibas. *Wireless Sensor Networks: An Information Processing Approach*. Morgan Kaufmann, July 6, 2004. ISBN 978-1558609143.
- [1836] Holger Karl and Andreas Willig. *Protocols and Architectures for Wireless Sensor Networks*. Wiley & Sons, first edition, June 24, 2005. ISBN 978-0470095102.
- [1837] Rajeev Shorey, A. Ananda, Mun Choon Chan, and Wei Tsang Ooi. *Mobile, Wireless, and Sensor Networks: Technology, Applications, and Future Directions*. Wiley-IEEE Press, March 23, 2006. ISBN 978-0471718161.
- [1838] Fred M. Discenzo, Dukki Chung, and Kenneth A. Loparo. Power Scavenging Enables Maintenance-Free Wireless Sensor Nodes. In *Proceedings of NECSI International Conference on Complex Systems*, June 25–30, 2006, New England Complex Systems Institute, Marriott Boston Quincy, Boston, MA, USA. Online available at <http://necsi.org/events/iccs6/viewpaper.php?id=188> [accessed 2007-08-01].
- [1839] Shad Roundy, Dan Steingart, Luc Frechette, Paul Wright, and Jan Rabaey. Power Sources for Wireless Sensor Networks. In *EWSN 2004: European workshop on wireless sensor networks No1*, volume 2920/2004 of *Lecture Notes in Computer Science (LNCS)*, pages 1–17. Springer Berlin / Heidelberg, 2004, Berlin, Germany. ISBN 978-3-540-20825-9. Online available at <http://www.eureka.gme.usherb.ca/memslab/docs/PowerReview-2.pdf> [accessed 2007-08-01].
- [1840] Joseph A. Paradiso and Thad Starner. Energy Scavenging for Mobile and Wireless Electronics. *IEEE Pervasive Computing*, 04(1):18–27, 2005. ISSN 1536-1268. doi:10.1109/MPRV.2005.9. Online available at <http://www.media.mit.edu/resenv/papers.html> and <http://whitepapers.silicon.com/0,39024759,60295509p,00.htm> [accessed 2007-08-01].
- [1841] Mohammad Rahimi, Hardik Shah, Gaurav Sukhatme, John Heidemann, and Deborah Estrin. Studying the Feasibility of Energy Harvesting in a Mobile Sensor Network. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 19–24, May 2003. IEEE, Taipei, Taiwan. Online available at <http://www>.

- isi.edu/~johnh/PAPERS/Rahimi03a.pdf and <http://www.citeulike.org/user/cr106/article/935917> [accessed 2007-08-01].
- [1842] Farhan Simjee and Pai H. Chou. Everlast: long-life, supercapacitor-operated wireless sensor node. In *ISLPED'06: Proceedings of the 2006 international symposium on Low power electronics and design*, pages 197–202, October 4–6, 2006, Tegernsee, Bavaria, Germany. ACM Press, New York, NY, USA. ISBN 1-59593-462-6. Online available at <http://portal.acm.org/citation.cfm?id=1098918.1098980> [accessed 2007-08-01].
 - [1843] Jaein Jeong, Xiaofan Fred Jiang, and David E. Culler. Design and Analysis of Micro-Solar Power Systems for Wireless Sensor Networks. Technical Report UCB/EECS-2007-24, EECS Department, University of California, Berkeley, February 2007. Online available at <http://www.eecs.berkeley.edu/Pubs/TechRpts/2007/EECS-2007-24.html> [accessed 2007-08-01].
 - [1844] Dusit Niyato, Ekram Hossain, and Afshin Fallahi. Sleep and Wakeup Strategies in Solar-Powered Wireless Sensor/Mesh Networks: Performance Analysis and Optimization. *IEEE Transactions on Mobile Computing*, 6(2):221–236, 2007. ISSN 1536-1233.
 - [1845] Francis R. Szabo and Paul E. Kladitis. Design, Modeling and Testing of Polysilicon Optothermal Actuators for Power Scavenging Wireless Microrobots. In *Proceedings of the 2004 International Conference on MEMS, NANO and Smart Systems, 2004. ICMENS 2004*, pages 446–452, August 25–27, 2004. IEEE Computer Society, Los Alamitos, CA, USA. ISBN 0-7695-2189-4.
 - [1846] Lei Wang and F. G. Yuan. Energy harvesting by magnetostrictive material (MsM) for powering wireless sensors in SHM. In *SPIE Smart Structures and Materials & NDE and Health Monitoring, 14th International Symposium (SSN07), 2007 SPIE/ASME Best Student Paper Presentation Contest*, March 18–22, 2007. Online available at <http://www.mae.ncsu.edu/research/SSML/paper.html> and <http://adsabs.harvard.edu/abs/2007SPIE.6529E.121W> [accessed 2007-08-01].
 - [1847] Joseph A. Paradiso. Systems for human-powered mobile computing. In *DAC'06: Proceedings of the 43rd annual conference on Design automation*, pages 645–650, 2006, San Francisco, CA, USA. ACM Press, New York, NY, USA. ISBN 1-59593-381-6. Online available at <http://www.media.mit.edu/resenv/pubs/papers/2006-07-DAC-Paper.pdf> and <http://portal.acm.org/citation.cfm?id=1147074> [accessed 2007-08-01].
 - [1848] John Kymissis, Clyde Kendall, Joseph Paradiso, and Neil Gershenfeld. Parasitic Power Harvesting in Shoes. In *ISWC'98: Proceedings of the 2nd IEEE International Symposium on Wearable Computers*, page 132, 1998. IEEE Computer Society, Washington, DC, USA. ISBN 0-8186-9074-7. Online available at <http://www.media.mit.edu/resenv/papers.html> and <http://citeseer.ist.psu.edu/kymissis98parasitic.html> [accessed 2007-08-01].
 - [1849] Katja Schwieger, Heinrich Nuszowski, and Gerhard Fettweis. Analysis of Node Energy Consumption in Sensor Networks. In Holger Karl, Andreas Willig, and Adam Wolisz, editors, *Proceedings of Wireless Sensor Networks, First European Workshop (EWSN)*, volume 2920 of *Lecture Notes in Computer Science (LNCS)*, pages 94–105. Springer, January 19–21, 2004, Berlin, Germany. ISBN 3-540-20825-9. Online available at <http://citeseer.ist.psu.edu/634903.html> [accessed 2007-08-01].
 - [1850] Eric Fleury, Jean-Loup Guillaume, Céline Robardet, and Antoine Scherrer. Analysis of Dynamic Sensor Networks: Power Law Then What? In *Second International Conference on COMMunication Systems softWARE and middlewaRE (COM-SWARE 2007)*. IEEE, January 7–12, 2007, Bangalore, India. ISBN 1-4244-0614-5. doi:10.1109/COMSWA.2007.382427. Online available at <http://jlguillaume.free.fr/www/publis/Guillaume07analysis.pdf> [accessed 2007-08-01].
 - [1851] Guillaume Chelius, Eric Fleury, and Thierry Mignon. Lower and Upper Bounds for Minimum Energy Broadcast and Sensing Problems in Sensor Networks. *International Journal of Parallel, Emergent and Distributed Systems*, to be published, 2005. Also:

- technical report, January 2004: Rapport de recherche de l'INRIA - Rhone-Alpes, Equipe: ARES. Online available at <http://www.inria.fr/rrrt/rr-5072.html> [accessed 2007-08-01]. See also [1961].
- [1852] Q. Gao, K. J. Blow, D. J. Holding, and I. Marshall. Analysis of Energy Conservation in Sensor Networks. *Wireless Networks*, 11(6):787–794, November 2005. ISSN 1022-0038 (Print) 1572-8196 (Online). Online available at <http://www.cs.kent.ac.uk/pubs/2005/2193/content.pdf> and <http://portal.acm.org/citation.cfm?id=1160416> [accessed 2007-08-01].
- [1853] Chee-Yee Chong and S.P. Kumar. Sensor networks: evolution, opportunities, and challenges. *Proceedings of the IEEE*, 91(8):1247–1256, August 2003. doi:10.1109/JPROC.2003.814918.
- [1854] H. Durrant-Whyte. Data Fusion in Sensor Networks. In *Proceedings of IEEE International Conference on Video and Signal Based Surveillance AVSS'06*, pages 39–39, November 2006, Sydney, Australia. ISBN 0-7803-9202-7. Online available at http://www.ee.ucla.edu/~mbs/ipsn05/keynote_abstracts/hdurrantwhyte.pdf and <http://portal.acm.org/citation.cfm?id=1147687> [accessed 2007-08-01].
- [1855] Ramesh Govindan, Joseph M. Hellerstein, Wei Hong, Samuel Madden, Michael Franklin, and Scott Shenker. The Sensor Network as a Database. Technical Report 02-771, University of South California, Computer Science Department, September 2002. Online available at <ftp://ftp.usc.edu/pub/csinfo/tech-reports/papers/02-771.pdf> and <http://citeseer.ist.psu.edu/601935.html> [accessed 2007-08-01].
- [1856] Jane Tateson, Christopher Roadknight, Antonio Gonzalez, Taimur Khan, Steve Fitz, Ian Henning, Nathan Boyd, Chris Vincent, and Ian Marshall. Real World Issues in Deploying a Wireless Sensor Network for Oceanography. In *Workshop on Real-World Wireless Sensor Networks REALWSN'05*, June 20–21, 2005. Stockholm, Sweden. Online available at <http://www.cs.kent.ac.uk/pubs/2005/2209/content.pdf> and <http://www.sics.se/realwsn05/papers/tateson05realworld.pdf> [accessed 2007-09-15].
- [1857] Bryan Horling, Roger Mailler, Mark Sims, and Victor Lesser. Using and Maintaining Organization in a Large-Scale Distributed Sensor Network. *Proceedings of the Workshop on Autonomy, Delegation, and Control (AAMAS03)*, July 2003. Online available at <ftp://mas.cs.umass.edu/pub/bhorling/03-organization.ps.gz> and <http://mas.cs.umass.edu/paper/247> [accessed 2007-08-01].
- [1858] Chun-Hsin Wu. Peer-to-peer Systems: Macro-computing with Micro-computers, July 2003. Presented at 3rd International Conference on Open Source in Taipei, Taiwan. Presentation available at <http://www.csie.nuk.edu.tw/~wuch/publications/2003-icos-p2p-wuch.pdf> [accessed 2007-08-13].
- [1859] David Culler and Dr. David Tennenhouse. Largest Tiny Network Yet – Large-Scale Demonstration of Self-Organizing Wireless Sensor Networks, August 2001, San Jose, California, USA. See <http://webs.cs.berkeley.edu/800demo/> [accessed 2007-07-03].
- [1860] Robert Szewczyk, Joe Polastre, Alan Mainwaring, John Anderson, and David Culler. An Analysis of a Large Scale Habitat Monitoring Application. In *Proceedings of The Second ACM Conference on Embedded Networked Sensor Systems SenSys 2004*, pages 214–226, November 3–5, 2004, Baltimore, MD. Online available at <http://www.eecs.harvard.edu/~mdw/course/cs263/papers/> and <http://portal.acm.org/citation.cfm?doid=1031495.1031521> [accessed 2007-08-01].
- [1861] Vinod Subramanian, Rajkumar Arumugam, and Ali A. Minai. Self-Organization of Connectivity and Geographical Routing in Large-Scale Sensor Networks. In Ali Minai, Dan Braha, Helen Harte, Larry Rudolph, Temple Smith, Gunter Wagner, and Yaneer Bar-Yam, editors, *Proceedings of the Fourth International Conference on Complex Systems*, June 9–14, 2002, Nashua, NH. Online available at http://necsi.org/events/iccs/2002/NAp07_subramanian_iccs4-1.pdf [accessed 2007-08-01].

- [1862] Tommaso Melodia, Dario Pompili, and Ian F. Akyildiz. On the interdependence of Distributed Topology Control and Geographical Routing in Ad Hoc and Sensor Networks. *Journal of Selected Areas in Communications*, 23(3):520–532, March 2005. doi:10.1109/JSAC.2004.842557.
- [1863] Jian Chen, Yong Guan, and Udo Pooch. Customizing a Geographical Routing Protocol for Wireless Sensor Networks. In *ITCC'05: Proceedings of the International Conference on Information Technology: Coding and Computing (ITCC'05) - Volume II*, volume 2, pages 586–591, April 4–6, 2005. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2315-3. doi:10.1109/ITCC.2005.118.
- [1864] M. Brian Blake. Coordinating multiple agents for workflow-oriented process orchestration. *Information Systems and E-Business Management*, 1(4):387–404, November 2003. ISSN 1617-9846 (Print) 1617-9854 (Online). doi:10.1007/s10257-003-0023-1. Online available at <http://www.springerlink.com/content/tk7jdgbraq0pywcr/fulltext.pdf> and http://wotan.liu.edu/docis/dbl/isebbm/2003_1_4_387_CMAFWP.htm [accessed 2007-08-13].
- [1865] Jesper M. Johansson, Salvatore T. March, and J. David Naumann. Modeling Network Latency and Parallel Processing in Distributed Database Design. *Decision Sciences*, 34(4):677–706, November 2003. Online available at http://findarticles.com/p/articles/mi_qa3713/is_200310/ai_n9253628 [accessed 2007-08-13].
- [1866] Robert Beverly, Karen Sollins, and Arthur Berger. SVM Learning of IP Address Structure for Latency Prediction. In *MineNet'06: Proceedings of the 2006 SIGCOMM workshop on Mining network data*, pages 299–304, September 2006, Pisa, Italy. ACM Press, New York, NY, USA. Online available at <http://www.sigcomm.org/sigcomm2006/papers/minenet-04.pdf> and <http://portal.acm.org/citation.cfm?id=1162678.1162682> [accessed 2007-08-13].
- [1867] James Z. Wang and Matti Vanninen. Self-Configuration Protocols for Small-Scale P2P Networks. In *Proceedings of 10th Network Operations and Management Symposium, NOMS 2006*, pages 1–4. IEEE/IFIP, April 2006, Vancouver, Canada. ISBN 1-4244-0142-9. Online available at <http://www.cs.clemson.edu/~jzwang/pub/nomshort.pdf> [accessed 2007-08-19].
- [1868] Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching Agreement in the Presence of Faults. *Journal of the ACM*, 27(2):228–234, 1980. ISSN 0004-5411. Online available at <http://research.microsoft.com/users/lamport/pubs/reaching.pdf> and <http://research.microsoft.com/users/lamport/pubs/reaching.pdf> [accessed 2007-08-13].
- [1869] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine Generals Problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982. ISSN 0164-0925. Online available at <http://research.microsoft.com/users/lamport/pubs/byz.pdf> and <http://portal.acm.org/citation.cfm?id=357176> [accessed 2007-08-13].
- [1870] Leslie Lamport. The Weak Byzantine Generals Problem. *Journal of the ACM*, 30(3):668–676, 1983. ISSN 0004-5411. Online available at <http://doi.acm.org/10.1145/2402.322398> and <http://research.microsoft.com/users/lamport/pubs/weak-byz.pdf> [accessed 2007-08-13].
- [1871] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance and proactive recovery. *ACM Trans. Comput. Syst.*, 20(4):398–461, 2002. ISSN 0734-2071. Online available at <http://research.microsoft.com/users/mcastro/publications.htm> and <http://portal.acm.org/citation.cfm?id=571640> [accessed 2007-08-13].
- [1872] Jean-Philippe Martin, Lorenzo Alvisi, and Michael Dahlin. Minimal Byzantine Storage. In *Distributed Computing, 16th international Conference, DISC 2002*, pages 311–325, October 2002. ISBN 3-540-00073-9. Online available at <http://www.cs.utexas.edu/users/jpmartin/papers/MinByz-TR.ps> and <http://www.springerlink.com/content/5ylhktru1bh994fv/fulltext.pdf> [accessed 2007-08-13].

- [1873] Jean-Philippe Martin, Lorenzo Alvisi, and Michael Dahlin. Small Byzantine Quorum Systems. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 374–383, June 2002. Online available at http://www.cs.utexas.edu/users/lorenzo/papers/smallByz_DSN.pdf and <http://citeseer.ist.psu.edu/593329.html> [accessed 2007-08-13].
- [1874] Jian Yin, Jean-Philippe Martin, Arun Venkataramani, Lorenzo Alvisi, and Mike Dahlin. Separating agreement from execution for byzantine fault tolerant services. In *Proceedings of the nineteenth ACM symposium on Operating systems principles*, pages 253–267. ACM Press, 2003, Bolton Landing, NY, USA. ISBN 1-58113-757-5. Online available at <http://portal.acm.org/citation.cfm?id=1165389.945470> and <http://www.cs.utexas.edu/users/lasr/papers/Yin03Separating.pdf> [accessed 2007-08-13].
- [1875] J-P. Martin and L. Alvisi. A Framework for Dynamic Byzantine Storage. In *Proceedings of the International Conference on Dependable Systems and Networks*, June 2004. Online available at <http://www.cs.utexas.edu/users/lasr/papers/Martin05Fast.pdf> [accessed 2007-08-13].
- [1876] J-P. Martin and L. Alvisi. Fast Byzantine Consensus. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 402–411, June 2005.
- [1877] Harry C. Li, Allen Clement, Edmund L. Wong, Jeff Napper, Indrajit Roy, Lorenzo Alvisi, and Michael Dahlin. BAR Gossip. In *Proceedings of the 2006 USENIX Operating Systems Design and Implementation (OSDI)*, November 2006. Online available at <http://www.cs.utexas.edu/users/dahlin/papers.html> [accessed 2007-08-13].
- [1878] Amitanand S. Aiyer, Lorenzo Alvisi, Allen Clement, Michael Dahlin, Jean-Philippe Martin, and Carl Porth. BAR Fault Tolerance for Cooperative Services. In *20th ACM Symposium on Operating Systems Principles*, October 2005. Online available at <http://www.cs.utexas.edu/users/lorenzo/papers/sosp05.pdf> and <http://doi.acm.org/10.1145/1095809.1095816> [accessed 2007-08-13].
- [1879] Artem V. Chebotko from the Wayne State University (WSU), Detroit, USA. *Programming Languages (Course Material CSC3200)*. Wayne State University (WSU), Detroit, USA, 2006. Online available at <http://www.cs.wayne.edu/~artem/> [accessed 2007-07-03].
- [1880] Peter Köchel from the TU Chemnitz, Fakultät für Informatik, Professur Modellierung und Simulation, Straße der Nationen 62, 09107 Chemnitz, Germany. *Algorithmen und Programmierung (Course Material)*. TU Chemnitz, Fakultät für Informatik, Professur Modellierung und Simulation, Straße der Nationen 62, 09107 Chemnitz, Germany, 2007. Online available at <http://www.tu-chemnitz.de/informatik/ModSim/> [accessed 2007-07-03].
- [1881] Noam Chomsky. *Syntactic structures*. 's-Gravenhage: Mouton & Co., 1957. Online available at http://books.google.de/books?id=a6a_b-CXYAkC and http://web.uni-marburg.de/dsa/Direktor/Rabanus/pdf/Syntactic_Structures.pdf [accessed 2007-09-14].
- [1882] Dick Grune and Criel J. H. Jacobs. *Parsing techniques a practical guide*. Ellis Horwood Limited, Chichester, England, August 1991. ISBN 978-0136514312. Online available at <http://citeseer.ist.psu.edu/grune90parsing.html> [accessed 2007-09-14].
- [1883] Noam Chomsky. Three models for the description of language. *IEEE Transactions on Information Theory*, 2(3):113–124, September 1956. Online available at <http://www.chomsky.info/articles/195609--.pdf> [accessed 2007-09-14].
- [1884] Noam Chomsky and Marcel P. Schützenberger. The algebraic theory of context-free languages. In P. Braffort and D. Hirschberg, editors, *Computer Programming and Formal Systems*, pages 118–161. North-Holland, Amsterdam, 1963.
- [1885] András Kornai. Natural languages and the Chomsky hierarchy. In *Proceedings of the second conference on European chapter of the Association for Computational*

- Linguistics*, pages 1–7, 1985, Geneva, Switzerland. Association for Computational Linguistics, Morristown, NJ, USA. doi:10.3115/976931.976932. Online available at <http://portal.acm.org/citation.cfm?doid=976931.976932> and <http://www.aclweb.org/anthology-new/E/E85/E85-1001.pdf> [accessed 2007-09-14].
- [1886] J. W. Backus, J. H. Wegstein, A. van Wijngaarden, M. Woodger, F. L. Bauer, J. Green, C. Katz, J. McCarthy, A. J. Perlis, H. Rutishauser, K. Samelson, and B. Vauquois. Report on the Algorithmic Language ALGOL 60. *Communications of the ACM*, 3(5):299–314, May 1960. ISSN 0001-0782. See also <http://www.masswerk.at/algol60/> [accessed 2007-09-15] and [1962].
- [1887] Donald E. Knuth. Backus Normal Form vs. Backus Naur Form. *Communications of the ACM*, 7(12):735–736, 1964. ISSN 0001-0782. doi:10.1145/355588.365140. Online available at <http://doi.acm.org/10.1145/355588.365140> [accessed 2007-09-15].
- [1888] International Organization for Standardization (ISO). *ISO/IEC 14977:1996: Information technology – Syntactic metalanguage – Extended BNF*. International Organization for Standardization (ISO), 1, ch. de la Voie-Creuse, Case postale 56, CH-1211 Geneva 20, Switzerland, 1996. Online available at <http://www.cl.cam.ac.uk/~mgk25/iso-14977.pdf> [accessed 2007-09-15].
- [1889] Richard E. Pattis. EBNF: A Notation to Describe Syntax, the late 1980s. Online available at <http://www.cs.cmu.edu/~pattis/misc/ebnf.pdf> [accessed 2007-07-03].
- [1890] Donald E. Knuth. Semantics of Context-Free Languages. *Theory of Computing Systems/Mathematical Systems Theory*, 2(2):127–145, 1968. ISSN 1432-4350 (Print) 1433-0490 (Online). doi:10.1007/BF01692511. See [1891]. Online available at <http://www.springerlink.com/content/m2501m07m4666813/fulltext.pdf> [accessed 2007-09-15].
- [1891] Donald E. Knuth. Correction: Semantics of Context-Free Languages. *Theory of Computing Systems/Mathematical Systems Theory*, 5(1):95–96, 1971. ISSN 1432-4350 (Print) 1433-0490 (Online). doi:10.1007/BF01702865. See [1890]. Online available at <http://www.springerlink.com/content/rj10u682v25g6506/fulltext.pdf> [accessed 2007-09-15].
- [1892] Donald E. Knuth. The genesis of attribute grammars. In *WAGA: Proceedings of the international conference on Attribute grammars and their applications*, pages 1–12, 1990, Paris, France. Springer-Verlag New York, Inc., New York, NY, USA. ISBN 0-387-53101-7. Online available at <http://www.dcs.warwick.ac.uk/~sk/cs325/gag.pdf> [accessed 2007-09-15].
- [1893] Jukka Paakki. Attribute grammar paradigms – a high-level methodology in language implementation. *ACM Computing Surveys (CSUR)*, 27(2):196–255, 1995. ISSN 0360-0300. doi:10.1145/210376.197409. Online available at <http://doi.acm.org/10.1145/210376.197409> [accessed 2007-09-15].
- [1894] Alex Aiken from the UC Berkely, Electrical Engineering and Computer Science, College of Engineering, Computer Science Division. Lecture Notes CS 264, Spring 1995. Online available at <http://www.cs.berkeley.edu/~aiken/cs264/lectures/attribute-grammars> [accessed 2007-07-03].
- [1895] Nelson Correa. An extension of earley’s algorithm for S-attributed grammars. In *Proceedings of the fifth conference on European chapter of the Association for Computational Linguistics*, pages 299–302, 1991, Berlin, Germany. Association for Computational Linguistics, Morristown, NJ, USA. doi:10.3115/977180.977232. Online available at <http://dx.doi.org/10.3115/977180.977232> and <http://www.aclweb.org/anthology-new/E/E91/E91-1052.pdf> [accessed 2007-09-115].
- [1896] David A. Watt. An extended attribute grammar for Pascal. *ACM SIGPLAN Notices*, 14(2):60–74, 1979. ISSN 0362-1340. doi:10.1145/954063.954071. Online available at <http://doi.acm.org/10.1145/954063.954071> [accessed 2007-09-15].
- [1897] Ole Lehrmann Madsen. On defining semantics by means of extended attribute grammars. In *Semantics-Directed Compiler Generation, Proceedings of a Workshop*, vol-

- ume 94 of *Lecture Notes In Computer Science (LNCS)*, pages 259–299, January 14–18, 1980, Aarhus, Denmark. Springer-Verlag, London, UK. ISBN 3-540-10250-7.
- [1898] Henning Christiansen. Programming as language development. Technical Report 15, Department of Computer Science, Roskilde University, 1988. Ph.D. thesis (summary).
 - [1899] Henning Christiansen. The syntax and semantics of extensible languages. Technical Report 14, Department of Computer Science, Roskilde University, 1988.
 - [1900] Henning Christiansen. Parsing and compilation of generative languages. Technical Report 3, Department of Computer Science, Roskilde University, 1986. abridged version: [1901].
 - [1901] Henning Christiansen. Recognition of generative languages. In *Proceedings of Programs as Data Objects Workshop*, volume 217 of *Lecture Notes in Computer Science (LNCS)*, pages 63–81. Springer-Verlag, October 1985, Copenhagen, Denmark. ISBN 0-387-16446-4. Abridged version of [1900].
 - [1902] Aravind K. Joshi. How Much Context-Sensitivity is Necessary for Characterizing Structural Descriptions – Tree Adjoining Grammars. In D. Dowty, L. Karttunen, and A. Zwicky, editors, *Natural Language Processing – Theoretical, Computational and Psychological Perspective*, pages 206–250. Cambridge University Press, New York, NY, 1985. Originally presented in 1983.
 - [1903] Owen Rambow and Aravind K. Joshi. A Formal Look at Dependency Grammars and Phrase-Structure Grammars, with Special Consideration of Word-Order Phenomena. *Meaning-Text Theory*, 1997. Online available at <http://arxiv.org/abs/cmp-lg/9410007v1> [accessed 2007-09-15].
 - [1904] Aravind K. Joshi and Owen Rambow. A Formalism for Dependency Grammar Based on Tree Adjoining Grammar. In *Proceedings of the Conference on Meaning-Text Theory, MTT 2003*, June 16–18, 2003, Paris, France. Online available at <http://www1.cs.columbia.edu/~rambow/papers/joshi-rambow-2003.pdf> [accessed 2007-09-15].
 - [1905] John L. McCarthy. Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I. *Communications of the ACM*, 3(4):184–195, April 1960. ISSN 0001-0782. Online available at <http://doi.acm.org/10.1145/367177.367199> and <http://citeseer.ist.psu.edu/mccarthy60recursive.html> [accessed 2007-09-15].
 - [1906] John L. McCarthy, Paul W. Abrahams, Daniel J. Edwards, Timothy P. Hari, and Michael L. Levin. *LISP 1.5 Programmer’s Manual*. The MIT Press, August 1962. ISBN 978-0262130110, 0-262-13011-4. Second Edition, 15th Printing, Online available at <http://www.softwarepreservation.org/projects/LISP/book/LISP%201.5%20Programmers%20Manual.pdf> [accessed 2007-09-15]. See also [1906].
 - [1907] John L. McCarthy. History of LISP. In Richard L. Wexelblat, editor, *History of Programming Languages: Proceedings of the ACM SIGPLAN Conference*, pages 173–197. Academic Press, June 1978. Online available at <http://citeseer.ist.psu.edu/mccarthy78history.html> [accessed 2007-09-15]. See also <http://www-formal.stanford.edu/jmc/history/lisp/lisp.html> [accessed 2007-09-15].
 - [1908] R. Kent Dybvig. *The SCHEME programming language*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1987. ISBN 0-13-791864-X, 978-0-262-54148-0. With Illustrations by Jean-Pierre Hébert. Third edition (October 2003) Online available at <http://www.scheme.com/tspl3/> [accessed 2007-09-15].
 - [1909] Ron Rivest. S-Expressions, May 4, 1997. draft-rivest-sexp-00.txt. Online available at <http://people.csail.mit.edu/rivest/Sexp.txt> [accessed 2007-07-03]. Network Working Group, Internet Draft, Expires November 4, 1997.
 - [1910] Thomas Weise, Michael Zapf, Mohammad Ullah Khan, and Kurt Geihs from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Genetic Programming meets Model-Driven Development. *Kasseler Informatikschriften (KIS)* 2007, 2, University of Kassel, July 2, 2007. Persistent Identifier: urn:nbn:de:hebis:34-2007070218786. Online available at <http://www.it-weise.de/documents/files/WZKG2007DGPFd>.

- pdf and <http://kobra.bibliothek.uni-kassel.de/handle/urn:nbn:de:hebis:34-2007070218786> [accessed 2007-09-17].
- [1911] Michael Zapf and Thomas Weise from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Offline Emergence Engineering For Agent Societies. In *Proceedings of the Fifth European Workshop on Multi-Agent Systems EUMAS'07*, December 14, 2007, Elmouradi Hotel, Hammamet, Tunisia. Online available at <http://www.it-weise.de/documents/files/ZW2007EUMASTR.pdf> [accessed 2008-08-12], see also [1963].
 - [1912] Branko Souček and The IRIS Group, editors. *Dynamic, Genetic, and Chaotic Programming: The Sixth-Generation*. Sixth Generation Computer Technologies. John Wiley Interscience, New York, April 1992. ISBN 047155717X, 978-0471557173. Partly online available at http://www.amazon.com/gp/reader/047155717X/ref=sib_dp_pt/002-6076954-4198445#reader-link [accessed 2008-05-29].
 - [1913] Dario Floreano, Jean-Daniel Nicoud, and Francesco Mondada, editors. *Advances in Artificial Life, Proceedings of the 5th European Conference, ECAL'99*, volume 1674 of *Lecture Notes in Computer Science (LNCS)*, September 13–17, 1999, Lausanne, Switzerland. Springer. ISBN 3-540-66452-1.
 - [1914] Christoph Adami, Richard K. Belew, Hiroaki Kitano, and Charles E. Taylor, editors. *Artificial Life VI: Proceedings of the sixth international conference on Artificial life*, volume 6 of *Complex Adaptive Systems*, June 1998, Madison, Wisconsin, United States. MIT Press, Cambridge, MA, USA. ISBN 0-262-51099-5, 978-0262510998.
 - [1915] Federico Morán, Alvaro Moreno, Juan J. Merelo Guervós, and Pablo Chacón, editors. *Advances in Artificial Life, Proceedings of the Third European Conference on Artificial Life*, volume 929 of *Lecture Notes in Computer Science (LNCS)*, June 4–6, 1995, Granada, Spain. Springer. ISBN 3-540-59496-5.
 - [1916] George C. Williams. Pleiotropy, Natural Selection, and the Evolution of Senescence. *SAGE KE, Science of Aging Knowledge Environment*, 2001(1), October 3, 2001. cp13. See also [184].
 - [1917] Michael J. Kearns, Yishay Mansour, Andrew Y. Ng, and Dana Ron. An Experimental and Theoretical Comparison of Model Selection Methods. *Machine Learning*, 27(1): 7–50, 1997. See also [201].
 - [1918] Hans-Georg Beyer. Towards a Theory of ‘Evolution Strategies’: Some Asymptotical Results from the $(1, +\lambda)$ -Theory. Technical Report SYS-5/92, Fachbereich Informatik, Universität Dortmund, P.O. Box 500500, 44227 Dortmund, Germany, December 1 1992. See also [215].
 - [1919] Xin Yao, editor. *Progress in Evolutionary Computation, Selected Papers of the AI'93 and AI'94 Workshops on Evolutionary Computation Melbourne, Victoria, Australia, November 16, 1993 and Armidale, NSW, Australia, November 21-22, 1994*, volume 956/1995 of *Lecture Notes in Artificial Intelligence, subseries of Lecture Notes in Computer Science (LNCS)*, May 1995. Springer, Berlin/Heidelberg, Germany. ISBN 3-540-60154-6, 978-3-540-60154-8. doi:10.1007/3-540-60154-6.
 - [1920] L. Darell Whitley from the Computer Science Department, Colorado State University, Fort Collins. A Genetic Algorithm Tutorial. Technical Report CS-93-103, Computer Science Department, Colorado State University, Fort Collins, March 10, 1993. Online available at <http://citeseer.ist.psu.edu/177719.html> [accessed 2007-11-29]. See also [645].
 - [1921] Mathieu S. Capcarrère, Alex Alves Freitas, Peter J. Bentley, Colin G. Johnson, and Jon Timmis, editors. *Advances in Artificial Life, Proceedings of the 8th European Conference, ECAL 2005*, volume 3630 of *Lecture Notes in Computer Science (LNCS) subseries Lecture Notes in Artificial Intelligence (LNAI)*, September 5–9, 2005, University of Kent, Canterbury, Kent (UK). Springer Verlag. ISBN 3-540-28848-1.
 - [1922] Francisco J. Varela and Paul Bourguine, editors. *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, Decem-

- ber 11–13, 1991, Paris, France. A Bradford book, The MIT Press. ISBN 0-262-72019-1. Published in 1992.
- [1923] N. S. Sridharan, editor. *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, volume 1, August 1989, Detroit, MI, USA. Morgan Kaufmann, San Francisco, CA, USA. ISBN 1-55860-094-9. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-89-VOL1/CONTENT/content.htm> [accessed 2008-04-01]. See also [1964].
- [1924] Christopher G. Langton, C.E. Taylor, D.J. Farmer, and S. Rasmussen, editors. *Artificial Life II: Proceedings of the Workshop on Artificial Life*, volume X of *Santa Fe Institute Studies in the Science of Complexity*, February 1990, Santa Fe, New Mexico, USA. Addison-Wesley. ISBN 0201525712, 978-0201525717. Published 1992, republished by Westview Press (March 28, 2003).
- [1925] Jean-Arcady Meyer and Stewart W. Wilson, editors. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, September 24–28, 1990, Paris, France. The MIT Press, Cambridge, MA, USA. ISBN 0-262-63138-5. Published in 1991.
- [1926] Francisco Fernandez de Vega. *Modelos de Programacion Genetica Paralela y Distribuida con aplicaciones a la Sintesis Logica en FPGAs*. PhD thesis, University of Extremadura, 2001. Version español. For english version see [815]. Online available at <http://cum.unex.es/profes/profes/fcofdez/escritorio/investigacion/pgp/thesis/phd.html> [accessed 2007-09-09].
- [1927] J. L. Denenbourg, S. Goss, G. Nicolis, H. Bersini, and R. Dagonnier, editors. *ECAL-93 Self organisation and life: from simple rules to global complexity, Proceedings of the Second European Conference on Artificial Life*, May 24–26, 1993, Bruxelles, Belgium. Free University of Brussels.
- [1928] *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97*, volume 1, August 23–29, 1997, Nagoya, Japan. Morgan Kaufmann, San Francisco, CA, USA. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-97-VOL1/CONTENT/content.htm> [accessed 2008-04-01]. See also [1929, 1965].
- [1929] *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97*, volume 2, August 23–29, 1997, Nagoya, Japan. Morgan Kaufmann, San Francisco, CA, USA. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-97-VOL2/CONTENT/content.htm> [accessed 2008-04-01]. See also [1928, 1965].
- [1930] Pattie Maes, Maja J. Mataric, Jean-Arcady Meyer, Jordan B. Pollack, and Stewart W. Wilson, editors. *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior: From animals to animats 4*, September 9–13, 1996, Cape Code, USA. MIT Press, Cambridge, MA, USA. ISBN 0-262-63178-4.
- [1931] Thomas D. Haynes, Roger L. Wainwright, and Sandip Sen from the Department of Mathematical & Computer Sciences, The University of Tulsa. Evolving Cooperating Strategies. In Victor Lesser, editor, *Proceedings of the first International Conference on Multiple Agent Systems*, page 450. AAAI Press/MIT Press, June 12–14, 1995, San Francisco, USA. ISBN 0-262-62102-9. 1 page poster, see also [879].
- [1932] *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, volume 1, August 20–25, 1995, Montréal, Québec, Canada. Morgan Kaufmann, San Francisco, CA, USA. ISBN 1-55860-363-8. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-95-VOLt.htm> [accessed 2008-04-01]. See also [1966, 1967, 1968].
- [1933] Christopher G. Langton, editor. *Artificial Life III: Proceedings of the Workshop on Artificial Life*, volume XVII of *Santa Fe Institute Studies in the Sciences of Complexity*, June 15–19, 1992, Santa Fe, New Mexico, USA. Addison Wesley Longman. ISBN 020162494X, 978-0201624946. Published in August 1993.
- [1934] Justinian P. Rosca. *Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications*. Technical Report 95.2, University of Rochester, National Resource Laboratory for the Study of Brain and Behavior, Morgan Kauf-

- mann, San Mateo, California, Rochsater, New York, USA, July 9, 1995, Tahoe City, California, USA. Held in conjunction with the twelfth International Conference on Machine Learning.
- [1935] J. David Schaffer and Darell Whitley, editors. *Proceedings of COGANN-92. International Workshop on Combinations of Genetic Algorithms and Neural Networks*, June 6, 1992, Baltimore, Maryland, USA. IEEE Computer Society Press. ISBN 0-8186-2787-5, 978-0818627873. doi:10.1109/COGANN.1992.273951. INSPEC Accession Number: 4487943, Cat. No.92TH0435-8.
- [1936] *Proceedings of BIONETICS 2007, 2nd International Conference on Bio-Inspired Models of Network, Information, and Computing Systems*, December 10, 2007, Radisson SAS Beke Hotel, 43. Terez krt., Budapest H-1067, Hungary. Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering (ICST), IEEE, ACM. ISBN 978-963-9799-05-9.
- [1937] Wolfgang Banzhaf, Thomas Christaller, Peter Dittrich, Jan T. Kim, and Jens Ziegler, editors. *Advances in Artificial Life, Proceedings of the 7th European Conference, ECAL 2003*, volume 2801 of *Lecture Notes in Computer Science (LNCS)*, September 14–17, 2003, Dortmund, Germany. Springer. ISBN 3-540-20057-6.
- [1938] Hai H. Dam, Hussein A. Abbass, and Chris Lokan. DXCS: an XCS system for distributed data mining. Technical Report TR-ALAR-200504002, The Artificial Life and Adaptive Robotics Laboratory, School of Information Technology and Electrical Engineering, University of New South Wales, Northcott Drive, Campbell, Canberra, ACT 2600 Australia, 2005. ALAR Technical Report Series. Online available at <http://www.itee.adfa.edu.au/~alar/techreps/200504002.pdf> [accessed 2007-09-12]. See also [1201].
- [1939] Jean-Arcady Meyer, A. Berthoz, D. Floreano, H. Roitblat, and Stewart W. Wilson, editors. *SAB2000: From Animals to Animats 6, Proceedings of the 6th International Conference on Simulation of Adaptive Behavior*, August 2000, Paris, France. MIT Press, Cambridge, MA, USA. ISBN 0-262-63200-4.
- [1940] John Henry Holland. Escaping brittleness: the possibilities of general-purpose learning algorithms applied to parallel rule-based systems. *Computation & Intelligence: Collected Readings*, pages 275–304, 1995. See also [1234].
- [1941] John Mylopoulos and Raymond Reiter, editors. *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, volume 2, August 24–30, 1991, Sydney, Australia. Morgan Kaufmann, San Francisco, CA, USA. ISBN 1-55860-160-0. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-91-VOL2/CONTENT/content.htm> [accessed 2008-04-01] See Also PROC1991IJCAIV1, PROC1991IJCAITNGOIS, PROC1991IJCAIFLAF. C.
- [1942] Xavier Llorà and Kumara Sastry. Fast rule matching for learning classifier systems via vector instructions. Technical Report 2006001, IlliGAL, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, January 2006. Online available at <http://www.illigal.uiuc.edu/pub/papers/IlliGALs/2006001.pdf> [accessed 2007-09-12]. See also [1266].
- [1943] Pier Luca Lanzi, Daniele Loiacono, Stewart W. Wilson, and David E. Goldberg. Generalization in the XCSF Classifier System: Analysis, Improvement, and Extension. Technical Report 2005012, IlliGAL, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, March 2005. Online available at <http://www.illigal.uiuc.edu/pub/papers/IlliGALs/2005012.pdf> [accessed 2007-09-12]. See also [1274].
- [1944] Chii-Ruey Hwang. Simulated annealing: Theory and applications. *Acta Applicandae Mathematicae: An International Survey Journal on Applying Mathematics and Mathematical Applications*, 12(1):108–111, May 1988. ISSN 0167-8019 (Print) 1572-9036 (Online). doi:10.1007/BF00047572. Academia Sinica. This is a review of [1304].

- Online available at <http://www.springerlink.com/content/t1n71m85j33556m4/> [accessed 2007-08-25].
- [1945] Lester Ingber. Simulated Annealing: Practice versus Theory. Technical Report, Lester Ingber Research, P.O.B. 857, McLean, VA 22101, 1993. Online available at <http://ideas.repec.org/p/lei/ingber/93sa.html> and http://www.ingber.com/asa93_sapvt.pdf [accessed 2007-08-18]. See also [1310].
- [1946] Marco Dorigo, Gianni Di Caro, and L. Gambardella. Ant Algorithms for Discrete Optimization. *Artificial Life*, 5(2):137–172, 1999. Online available at http://www.idsia.ch/~luca/ij_23-alife99.pdf [accessed 2007-09-15]. See [1393].
- [1947] Edward Osborne Wilson. *Sociobiology: The New Synthesis*. Belknap Press, twenty-fifth anniversary edition edition, March 2000. ISBN 0674002350, 978-0674002357. See also [1398].
- [1948] De-Shuang Huang, Laurent Heutte, and Marco Loog, editors. *Advanced Intelligent Computing Theories and Applications. With Aspects of Contemporary Intelligent Computing Techniques – Proceedings of the Third International Conference on Intelligent Computing, ICIC 2007*, volume 2 of *Communications in Computer and Information Science*, August 21–24, 2007, Qingdao, China. Springer Berlin/Heidelberg. ISBN 978-3-540-74281-4 (Print) 978-3-540-74282-1 (Online). doi:10.1007/978-3-540-74282-1. Online available at <http://www.springerlink.com/content/k74022/> [accessed 2008-03-28].
- [1949] Jonathan E. Fieldsend, Richard M. Everson, and Sameer Singh. Extensions to the strength pareto evolutionary algorithm. *IEEE Transactions on Evolutionary Computation*, 2001. submitted, superseded by [1500]. Online available at <http://www.dcs.ex.ac.uk/people/reverson/pubs/espea.ps.gz> and <http://citeseer.ist.psu.edu/fieldsend01extensions.html> [accessed 2007-08-25].
- [1950] P. N. Suganthan, N. Hansen, J. J. Liang, Kalyanmoy Deb, Y. P. Chen, Anne Auger, and S. Tiwari. Problem Definitions and Evaluation Criteria for the CEC 2005 Special Session on Real-Parameter Optimization. Technical Report, Nanyang Technological University, Singapore, May 2005. Online available at http://www.ntu.edu.sg/home/epnsugan/index_files/CEC-05/Tech-Report-May-30-05.pdf [accessed 2007-10-07]. See also [1511, 496].
- [1951] Terry Jones. A Description of Holland’s Royal Road Function. Working Papers 94-11-059, Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA, November 1994. See also [1532].
- [1952] Pedro Lima, editor. *Robotic Soccer*. I-Tech Education and Publishing, Vienna, Austria, December 2007. ISBN 978-3-902613-21-9. Online available at <http://s.i-techonline.com/Book/Robotic-Soccer/ISBN978-3-902613-21-9.html> [accessed 2008-04-23].
- [1953] Bernhard Nebel, editor. *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, IJCAI 2001*, volume 1 and 2, August 4–10, 2001, Seattle, Washington, USA. Morgan Kaufmann, San Francisco, CA, USA. ISBN 1-55860-777-3. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-2001/content/content.htm> [accessed 2008-04-01]. See also [1969].
- [1954] Yoav Freund and Robert E. Schapire from the AT&T Labs, 180 Park Avenue, Florham Park, New Jersey 07932, USA. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55:119–139, 1997. Article no. SS971504. Online available at http://www.face-rec.org/algorithms/Boosting-Ensemble/decision-theoretic_generalization.pdf [accessed 2007-09-15]. See also [1582].
- [1955] Peter Bartlett, Yoav Freund, Wee Sun Lee, and Robert E. Schapire. Boosting the margin: a new explanation for the effectiveness of voting methods. *Annals of Statistics*, 26(5):1651–1686, 1998. Online available at <http://cs.nyu.edu/~cil217/ML/>

- BoostingMargin.pdf and <http://dx.doi.org/10.1214/aos/1024691352> [accessed 2007-09-15]. See also [1585].
- [1956] *Proceedings of 2006 IEEE Joint Conference on E-Commerce Technology (CEC'06) and Enterprise Computing, E-Commerce and E-Services (EEE'06)*, June 26–29, 2006, The Westin San Francisco Airport, 1 Old Bayshore Highway, Millbrae, United States. IEEE Computer Society, Los Alamitos, California, Washington, Tokyo. ISBN 978-0-7695-2511-2.
- [1957] *Proceedings of IEEE Joint Conference (CEC/EEE 2007) on E-Commerce Technology (9th CEC'07) and Enterprise Computing, E-Commerce and E-Services (4th EEE'07)*, July 23–26, 2007, National Center of Sciences, Tokyo, Japan. IEEE Computer Society, IEEE Computer Society. ISBN 978-0-7695-2913-4.
- [1958] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH'87, Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34, July 27–31, 1987, Anaheim, California, USA. Association for Computing Machinery (ACM), New York, NY, USA. ISBN 0-89791-227-6. doi:10.1145/37401.37406. Online available at <http://citeseer.ist.psu.edu/reynolds-flocks.html> and <http://doi.acm.org/10.1145/37401.37406> [accessed 2008-07-27]. See also [1675].
- [1959] Georg Gottlob and Toby Walsh, editors. *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, August 9–15, 2003, Acapulco, México. Morgan Kaufmann, San Francisco, CA, USA. ISBN 0-127-05661-0. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-2003/content.htm> [accessed 2008-04-01]. See also [1970].
- [1960] Steven R. Finch. *Mathematical Constants (Encyclopedia of Mathematics and its Applications)*. Cambridge University Press, August 18, 2003. ISBN 978-0521818056. See <http://algo.inria.fr/bsolve/> [accessed 2007-09-15].
- [1961] Guillaume Chelius, Eric Fleury, and Thierry Mignon. Lower and Upper Bounds for Minimum Energy Broadcast and Sensing Problems in Sensor Networks. In *ICPADS'05: Proceedings of the 11th International Conference on Parallel and Distributed Systems - Workshops (ICPADS'05)*, pages 88–92, 2005. IEEE Computer Society, Washington, DC, USA. ISBN 0-7695-2281-5. doi:10.1109/ICPADS.2005.181. See also [1851].
- [1962] J. W. Backus, F. L. Bauer, J. Green, C. Katz, J. McCarthy, A. J. Perlis, H. Rutishauser, K. Samelson, B. Vauquois, J. H. Wegstein, A. van Wijngaarden, and M. Woodger. Revised report on the algorithm language ALGOL 60. *Communications of the ACM*, 6(1):1–17, 1963. ISSN 0001-0782. doi:10.1145/366193.366201. Online available at <http://doi.acm.org/10.1145/366193.366201> and <http://www.masswerk.at/algol60/report.htm> [accessed 2007-09-15]. See also [1886].
- [1963] Michael Zapf and Thomas Weise from the University of Kassel, FB-16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany. Offline Emergence Engineering For Agent Societies. Kasseler Informatikschriften (KIS) 2007, 8, University of Kassel, FB16, Distributed Systems Group, Wilhelmshöher Allee 73, 34121 Kassel, Germany, December 7, 2007. Persistent Identifier: urn:nbn:de:hebis:34-2007120719844. Online available at <https://kobra.bibliothek.uni-kassel.de/handle/urn:nbn:de:hebis:34-2007120719844> and <http://www.it-weise.de/documents/files/ZW2007EUMASTR.pdf> [accessed 2007-11-20], see also [1911].
- [1964] N. S. Sridharan, editor. *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, volume 2, August 1989, Detroit, MI, USA. Morgan Kaufmann, San Francisco, CA, USA. ISBN 1-55860-094-9. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-89-VOL2/CONTENT/content.htm> [accessed 2008-04-01]. See also [1923].

- [1965] Anca L. Ralescu and James G. Shanahan, editors. *Fuzzy Logic in Artificial Intelligence, IJCAI'97 Selected and Invited Workshop Papers*, volume 1566 of *Lecture Notes in Computer Science (LNCS)*, 1999, Nagoya, Japan. Springer. ISBN 3-540-66374-6. See also [1928, 1929].
- [1966] *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, volume 2, August 20–25, 1995, Montréal, Québec, Canada. Morgan Kaufmann, San Francisco, CA, USA. ISBN 1-55860-363-8. Online available at <http://dli.iiit.ac.in/ijcai/IJCAI-95-VOL2/CONTENT/content.htm> [accessed 2008-04-01]. See also [1932, 1967, 1968].
- [1967] Gerhard Weiß and Sandip Sen, editors. *Adaption and Learning in Multi-Agent Systems, IJCAI'95 Workshop Proceedings*, volume 1042 of *Lecture Notes in Computer Science (LNCS)*, 1996, Montréal, Québec, Canada. Springer. ISBN 3-540-60923-7. See also [1932, 1966, 1968].
- [1968] Trevor P. Martin and Anca L. Ralescu, editors. *Fuzzy Logic in Artificial Intelligence, Towards Intelligent Systems, IJCAI'95 Selected Papers from Workshop*, volume 1188 of *Lecture Notes in Computer Science (LNCS)*, 1997, Montréal, Québec, Canada. Springer. ISBN 3-540-62474-0. See also [1932, 1966, 1967].
- [1969] Alexander Maedche, Steffen Staab, Claire Nedellec, and Eduard H. Hovy, editors. *IJCAI'2001 Workshop on Ontology Learning, Proceedings of the Second Workshop on Ontology Learning OL'2001*, volume 38 of *CEUR Workshop Proceedings*, August 4, 2001, Seattle, Washington, USA. CEUR-WS.org. Online available at <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-47/> [accessed 2008-04-01]. See also [1953]. Held in conjunction with the 17th International Conference on Artificial Intelligence IJCAI'2001.
- [1970] Subbarao Kambhampati and Craig A. Knoblock, editors. *Proceedings of IJCAI-03 Workshop on Information Integration on the Web (IIWeb-03)*, August 9–10, 2003, Acapulco, México. Online available at <http://www.isi.edu/info-agents/workshops/ijcai03/proceedings.htm> [accessed 2008-04-01]. See also [1959].

Index

- $(1 + 1) - ES$, 192
- $(GE)^2$, 152
- $(\mu', \lambda'(\mu, \lambda)^\gamma)$ -ES, 193
- $(\mu + 1)$ -ES, 192
- $(\mu + \lambda)$, 195
- $(\mu + \lambda)$ -ES, 192
- (μ, λ) -ES, 193
- $(\mu/\rho + \lambda)$ -ES, 193
- $(\mu/\rho, \lambda)$ -ES, 193
- χ^2 Distribution, 441
- $\frac{1}{5}$ -rule, 193
- $(\mu + \lambda)$, 74
- (μ, λ) , 74
- σ -algebra, 419
- Γ , 466
- Θ notation, 483
- 802.11, 363

- A* Search, 252
- AAAI, 61
- AAB, 359
- AAS, 359
- Abstraction, 481
- ACO, 233
- Action, 202
- AdaBoost, 327, 363
- Adapter, 365
- Adaptive Grammar, 508
- Adaptive Walk, 252
 - fitter dynamics, 253
 - greedy dynamics, 253
 - one-mutant, 252
- ADDO, 337

- Adenine, 24
- ADF, 135, 136, 162, 363
- ADG, 165
- Adjacency, 25, 26
- adjacent neighbors, 282
- Adjunction, 510, 511
- ADL, 164
- ADM, 136
- Admissible, 252
- AG, 505, 506
- AGA, 266
- Agent, 364
 - ant, 356
 - genetic, 357
- Agents, 77, 128, 195, 210
- Aggregate, 367
- Aggregate Function, 367
- Aggregation, 366, 367
 - gossip-based, 367
 - proactive, 367
 - reactive, 367
- Aggregation Protocols, 366, 367
- AI, 317
- AL, 178
- Algorithm, 479
 - abstraction, 481
 - anytime, 185
 - complexity, 482
 - determined, 482
 - determinism, 482
 - deterministic, 484
 - discrete, 481
 - distributed, 485

- euclidean, 305
- evaluate, 184
- evolutionary, 67
- evolve, 184
- finite, 482
- generational, 74
- Las Vegas, 484
- Monte Carlo, 484
- optimization, 30
- probabilistic, 4, 484
- randomized, 484
- termination, 482
- Algorithmic Chemistry, 169
- ALife, 178
- All-Or-Nothing, 186, 361
- Allele, 24
- Alphabet, 502
- ANN, 163, 318
- Ant
 - artificial, 9
 - forward, 356
- Ant Colony Optimization, 233
- Antisymmetrie, 415
- AntNet, 356, 357
- Anytime Algorithm, 185
- appleJuice, 490
- Application Server, 489
- Applications, 271
- Architecture
 - service oriented, 327
- Artificial Ant, 9, 303
- Artificial Chemistry, 178
- Artificial Embryogeny, 121
- Artificial Life, 178
- Asexual Reproduction, 117
- Assignment
 - soft, 178
- Asymmetrie, 415
- Athena, 364
- Attribute, 505
 - inherited, 505
 - synthesized, 505
- Attribute Grammar, 505
 - extended, 507, 508
 - L-attributed, 506
 - reflective, 153
 - S-attributed, 507
- Autoconstructive Evolution, 180
- Autocorrelation, 43
- Automatically Defined Functions, 135, 162
- Automatically Defined Groups, 165
- Automatically Defined Link, 164
- Automaton
 - cellular, 128, 195
- Auxiliary Tree, 510
- Average, 368
- Backus-Naur Form, 504
 - extended, 505
- Battery, 491
- Bayes Classifier
 - naïve, 318
- BBH, 124
- BDP, 260
- Bee, 199
- Bernoulli
 - distribution, 434
 - experiment, 434
 - trial, 434
- Best-First Search, 251
- BFS, 247
- BGP, 140
- Bias, 450
- Bibtex, 543
- Big-**O** notation, 482
- Big- **Ω** notation, 483
- Bijjective, 413
- Bijjectivity, 413
- BinInt, 38, 39, 289, 290
- Binomial Distribution, 434
- Biochemistry, 61, 128, 195, 224, 230, 242
- Biology, 61, 77, 215
- BIOMA, 77, 234
- BIONETS, 362
- Bird, 199
- BitCount, 289
- Bittorrent, 490
- Bloat, 45
- Block
 - building, 124
- BLUE, 454
- Bluetooth, 491
- BNF, 504
- Boosting, 327
- Bottleneck, 486
- Box-Muller, 462
 - polar, 463
- BPEL, 339
- BPEL4WS, 338
- Breadth-First Search, 247
- Broadcast, 498
- Broadcast-Distributed Parallel Evolutionary
 - Algorithm, 260
- BTNodes, 492
- Bucket Brigade, 204
- Building Block, 124
- Building Block Hypothesis, 124, 285

- Bus, 488
- Byzantine Fault, 501
- Candidate
 - solution, 23
- Cartesian Genetic Programming, 46, 165, 167
 - embedded, 167
- Catastrophe
 - complexity, 283
- Causality, 40, 59, 294
- CC, 359
- CDF, 421
 - continous, 422
 - discrete, 422
- CEC, 77, 234
- Cellular Automaton, 128, 195
- Cellular Encoding, 143
- Central Point Of Failure, 486
- Centroid, 469
- CFG, 503, 504
- CGE, 154
- CGP, 46, 165–167
 - embedded, 167
- Change
 - non-synonymous, 45
 - synonymous, 45
- Character String, 502
- Chemical Engineering, 61, 77, 114, 191, 194, 195, 219, 224, 239
- Chemistry, 61, 77, 114, 191, 194, 195, 219, 224, 239
 - algorithmic, 169
 - artificial, 178
- Chi-square Distribution, 441
- Chomsky Hierarchy, 503
- Christiansen
 - grammar, 154, 509
- Christiansen Grammar, 154
 - evolution, 154
- Christiansen Grammars, 509
- Chromosome, 117
- Chromosomes
 - string
 - fixed-length, 118
 - variable-length, 120
 - tree, 130
- CI, 80, 454
- Class
 - equivalence, 416
- Classifier, 202
- Classifier Systems, 197, 198, 322, 397
 - learning, 197, 204, 318
 - non-learning, 203
- Client, 488
- Client-Server, 257, 488
- Closure, 146, 188
- Clustering, 467
 - k -means, 472
 - n^{th} nearest neighbor, 472
 - algorithm, 468
 - hierarchical, 467
 - leader, 475
 - linkage, 473
 - partitional, 468, 471
 - partitions, 468
 - square error, 471
- Code Bloat, 343
- Codons, 140
- Coefficient
 - negative slope, 42
- Coefficient of Variation, 425
- Combination, 418
- Combinatorics, 418
- Communication, 61, 77, 114, 128, 210, 219, 230, 234, 242, 247
 - collective, 359
 - inter-process, 365
- Completeness, 25, 246
 - weak, 25
- Complexity Catastrophe, 283
- Compress, 167
- Computational Embryogeny, 121
- Computational Intelligence, 80
- Computer
 - science
 - theoretical, 479
- Computing
 - amorphous, 365
 - high-performance, 359
 - ubiquitous, 365
- Concatenation, 502
- Condition, 200
- Confidence
 - coefficient, 455
 - interval, 454
- Configuration, 364
- Connection Register, 170
- Content Sharing, 490
- Contest, 317
- Continuous Distributions, 435
- Contravariance, 330
- Convergence, 37
 - domino, 38, 289
 - premature, 37
 - prevention, 107
- Correlation

- fitness distance, 42
- genotype-fitness, 42
- operator, 42
- COST
 - action 239, 349
- Count, 423
- Covariance, 330, 426
- CPU, 359
- Creation, 108, 118, 120, 130, 163
- Credit Assignment Problem, 204
- Criterion
 - termination, 35
- Crossbow, 492
- Crossover, 70, 108, 119–121, 132
 - homologous, 161, 362
 - path, 354
 - point, 119
 - SAAN, 163
 - single-point, 119, 132
 - SSAAN, 164
 - SSIAN, 164
 - sticky, 161
 - strong context preserving, 132
 - tree, 131
- CS, 197
- CSG, 503
- CSP, 61, 77
- Cumulative Distribution Function, 421
- Cytosine, 24
- Dagstuhl Seminar, 78
- Data Mining, 77, 114, 128, 142, 191, 195, 197, 210, 224, 317, 467
- DATA-MINING-CUP, 317, 318
- Database Server, 489
- DE, 193, 194, 226
- Deceptiveness, 43, 48, 285
- Deceptivity, 43, 48
- Decile, 428
- Decision Maker
 - external, 18
- Decision Tree, 318
- Decreasing, 415
 - monotonically, 415
- Default Hierarchy, 201, 322
- Defense, 77
- Defined Length, 122
- DELB, 194
- Deme, 257
- Density Estimation, 457
 - crowding distance, 457
 - Kernel, 458
 - nearest neighbor, 457
 - Parzen window, 458
- Deoxyribonucleic acid, 23
- Deoxyribose, 24
- Depth-First Search, 248
 - iterative deepening, 249, 363
- Depth-limited Search, 249
- Derivation Tree, 503
- DERL, 194
- DES, 193
- Design, 61, 77, 194, 242
 - circuit, 77, 114, 128, 142, 167, 194, 195, 219
- Detector, 198
- Determined, 482
- Determinism, 482
- DFS, 248
- Differential Evolution, 193, 194, 226
- Differential Evolution Strategy, 193
- Dijkstra
 - 's algorithm, 353
- Discrete, 481
- Discrete Distributions, 429
- Distance
 - Euclidian, 470
 - Hamming, 469
 - Manhattan, 469
 - Measure, 469
- Distributed algorithms, 485
- Distribution, 255, 421, 429, 435
 - χ^2 , 441
 - chi-square, 441
 - continuous, 435
 - discrete, 429
 - exponential, 440, 463, 464
 - normal, 437, 462, 463
 - multivariate, 439
 - standard, 437
 - Poisson, 431
 - Student's t, 445
 - t, 445
 - uniform, 430, 436, 460, 462, 463
 - continuous, 436
 - discrete, 430
- Distribution Binomial, 434
- Diversity, 39, 188
- DMC, 317
- DNA, 23, 24, 140, 161
- do not Care, 322, 326
- Domination, 13
- Domino
 - convergence, 38, 289
- Don't Care, 122, 200
- Downhill Simplex, 223
- Drunkyard's Walk, 250
- Duplication, 108

- Dust Networks, 492
- EA, 67, 73, 76, 77, 80, 81, 366
- EA/AE, 78
- EAG, 507, 508
- EARL, 197
- EBNF, 505
- ECGP, 167
- ECJ, 154
- Economics, 61, 77, 114, 128, 152, 219
- Edge Encoding, 142
- Editing, 133
- EDL, 510
- Effector, 198
- Efficiency
 - Pareto, 13
- home, 490
- Elitism, 75
- Embrogeny, 121
 - artificial, 121
- Embryogenesis, 121
- Embryogenic, 121
- Embryogeny
 - artificial, 121
 - computational, 121
- EMO, 78
- EMOO, 68, 80
- Encapsulation, 134
- Encoding
 - cellular, 143
 - edge, 142
- Endnote, 543
- Energy Source, 491
- Engineering, 61, 77, 194, 242
 - electrical, 77, 114, 128, 142, 167, 194, 195, 219
- Entropy, 429
 - continuous, 429
 - differential, 429
 - information, 429
- Entscheidungsproblem, 184
- Environment, 198
 - protection, 77, 191
 - surveillance, 77, 191
- Eoarchean, 69
- EON, 349
- EP, 72, 195, 196
- Ephemeral Random Constants, 342
- Epistacy, 47
- Epistasis, 42, 47, 293, 299, 301, 302
 - in Genetic Programming, 168
 - in GPMs, 168
 - positional, 168
- Epistatic Road, 287
- Epistatic Variance, 42
- Equilibrium
 - punctuated, 44
- Equivalence
 - class, 416
 - relation, 416
- eRBGP, 176, 177
- ERL, 197
- Error, 450
 - burst, 499
 - mean square, 450
 - threshold, 42, 290
- Error Threshold, 42, 290
- ES, 72, 191, 192
- Estimation Theory, 450
- Estimator, 450
 - best linear unbiased, 454
 - maximum likelihood, 453
 - point, 450
 - unbiased, 450
- Euclidean Algorithm, 305
- Euclidian Distance, 470
- EUROGEN, 78, 115, 192, 196
- EuroGP, 128
- Evaluation, 34
- Event
 - certain, 417
 - conflicting, 417
 - elementary, 417
 - impossible, 417
 - random, 417
- EvoCOP, 78
- Evolution
 - autoconstructive, 180
 - differential, 193
- Evolution Strategy, 72, 191, 192
- Evolutionary Algorithm, 67, 76, 77, 80, 81
 - basic, 71
 - broadcast-distributed parallel, 260
 - cycle, 68
 - generational, 74
 - multi-objective, 68
 - parallelization, 256
 - steady state, 74
- Evolutionary Operation, 73, 223
 - randomized, 73
- Evolutionary Programming, 72, 166, 195, 196
- Evolutionary Reinforcement Learning, 197
- Evolvability, 42, 44
- EVOP, 73
- EvoWorkshops, 78
- Expand, 167
- expand, 245

- Expected value, 424
- Experiment
 - factorial, 273
- Exploitation, 39, 40
- Exploration, 39, 245
- Exponential Distribution, 440
- Extended Backus-Naur Form, 505
- External Decision Maker, 18
- Extinctive Selection, 73
 - left, 74
 - right, 74
- Extrema Selection, 46
- Factorial, 418
- FDC, 42
- FDL, 527
- FEA, 79
- File Sharing, 490
- Finance, 61, 77, 114, 128, 152, 219
- Finite, 482
- Finite State Machine, 126, 195, 304, 359, 361
- Fitness, 28
 - nature, 71
 - optimization, 71
- Fitness Assignment, 82
 - Pareto ranking, 83
 - Prevalence ranking, 83
 - Tournament, 92
 - weighted sum, 83
- Fitness Landscape, 29
 - deceptive, 43
 - ND, 285
 - neutral, 43
 - NK, 281
 - NKp, 284
 - NKq, 284
 - p-Spin, 284
 - rugged, 40
 - technological, 284
- Fly, 198
- FOGA, 115
- home, 490
- Forma, 42, 56, 57
 - analysis, 56
- Formae, 57
- Formal Grammar, 503
- Fraglets, 181, 362
 - evolution of, 361
- Free Lunch
 - no, 54
- Frequency
 - absolute, 419
 - relative, 419
- Frog, 198
- FSM, 128, 195, 359, 361
- Full, 130
- Fully Connected, 488
- Function, 414
 - ADF, 135, 162
 - aggregate, 367
 - automatically defined, 135, 162
 - benchmark, 279
 - cumulative distribution, 421
 - gamma, 466
 - monotone, 414
 - objective, 3
 - probability density, 422
 - probability mass, 422
 - synthesis, 128, 142, 159, 167
 - trap, 43, 285
- Function Set
 - relational, 350
- Functional, 413
- Functionality, 413
- FWGA, 115
- GA, 72, 113–116
- Gads, 147–149, 153, 168
 - 1, 147
 - 2, 153
- GAGS, 148
- GALESIA, 115
- Game, 77, 128, 195, 210
- Gamma, 466
- GARA, 354, 355
- Gauss-Markov Theorem, 454
- GCD, 305, 306
 - problem, 305
- GCL, 172
- GE, 149, 150, 152, 153, 168
- GECCO, 79, 115, 129, 234, 239
- Gene, 24
- Gene Expression Programming, 140, 142
- Generality, 50
- Generation, 34
- Generational, 74
- Generative Grammar, 502
- Genetic Algorithm, 72, 113–116, 206, 227
 - cellular, 261
 - cycle, 113
 - for deriving software, 147
 - grammar-based, 147
- Genetic Algorithms, 126
 - natural representation, 117
 - real-encoded, 117
- Genetic Network Programming, 164
- Genetic Programming, 72, 125, 127–130
 - binary, 140

- connective, 351
- crossover
 - homologous, 161
- decision trees, 351
- epistasis, 168
- grammar-guided, 145
- linear, 159, 160
 - page-based, 161
- node-pair encoding, 352
- ontogenic, 179
- parallel distributed, 162
- rule-based, 171
 - extended, 176
- standard, 126
- strongly typed, 146
- TAG, 156, 159
- tree-adjoining grammar-guided, 156, 159
- tree-based, 126, 127, 130
- Genetic Programming Kernel, 147, 360
- GENITOR, 74
- Genome, 23, 116
- Genomes
 - string, 118
 - tree, 130
- Genotype, 23
- Genotype-Phenotype Mapping, 25, 59, 121, 138
- Genotype-Phenotype mapping, 121
- Geometry, 114, 128, 142, 191, 219, 239
- GEP, 140–142
- GEWS, 152
- GFC, 42
- GGGP, 145
- Glass
 - spin, 29, 284
- Global Optimization Algorithm, 30
- GlobalAnt, 357
- GNP, 164, 165
- Gnutella, 353, 490
- Goal Attainment, 18
- Goal Programming, 18
- GP, 72, 125, 127–130
 - epistasis, 168
- GPK, 147, 360
- GPM, 25, 121, 138, 141
 - epistasis, 168
- GPTP, 129
- Gradient, 34
 - descend, 34
- Gradient Descent
 - stochastic, 212
- Grammar, 145
 - adaptive, 508
 - recursive, 509
- attribute, 505
- BNF, 504
- Christiansen, 509
 - evolution, 154
- context-free, 503, 504
- context-sensitive, 503
- derivation tree, 503
- EBNF, 505
- formal, 503
- generative, 502, 509
- induction, 128, 197
- L-attributed, 506
- recursive enumerable, 503
- regular, 503
- S-attributed, 507
- TAG, 510
- tree-adjoining, 510
 - lexicalized, 511
- tree-adjunct, 510
 - lexicalized, 511
- Grammatical Evolution, 149, 152, 153
 - Christiansen, 154, 168
- Granularity, 398
- Graph
 - butterfly, 372
- Greatest Common Divisor, 305
- Greedy Search, 251
- Grid, 488
- Grow, 130
- Guanine, 24
- Guarded Command Language, 172
- Halting Criterion, 35
- Halting Problem, 184, 186
 - reductio ad absurdum, 185
- Hamming Distance, 469
- HBGA, 113
- HC, 209, 210
- Herman, 125
- Heuristic, 4, 251
 - admissible, 252
 - monotonic, 252
- Heuristic Random Optimization, 214
- Hierarchy, 488
 - Chomsky, 503
 - default, 201, 322
- HiGP, 179
- Hill Climbing, 209, 210
 - multi-objective, 210
 - randomized restarts, 212
 - stochastic, 212
- HIS, 62
- Histogram, 457

- Homologous Crossover, 120, 161
- Homology, 161
- HRO, 214
- HSMP SO, 226
- HTTP, 364, 489
- Hydrogen Bond, 24
- Hyperplane, 122
- Hypothesis
 - building block, 124, 285
- IAAI, 62
- ICANNGA, 79, 115, 129
- ICGA, 115
- ICNC, 62, 79, 234, 239
- IDDFS, 249, 250, 331, 363
- IDS, 363
- IF-FOOD-AHEAD, 305
- Image Processing, 77, 114, 191, 197, 219, 467
- Implicit Parallelism, 70
- Increasing, 414
 - monotonically, 414
- Individual, 28
- IndividualPrinterPipe, 403
- inequalities
 - method of, 16
- Information
 - management, 467
 - processing, 467
- Information Management, 467
- Information Processing, 467
- Informed Search, 251
- Injective, 413
- Injectivity, 413
- Input, 481
- Input-Processing-Output, 125, 481
- Instant Messaging, 490
- Intelligence
 - artificial, 317
- Interval, 408
 - confidence, 454
- Intrinsic Parallelism, 70
- Intron, 117, 149, 162, 187
- Intrusion
 - detection, 363
- IPC, 365
- IPO, 125
- IPO Model, 481
- Irreflexivenss, 415
- isGoal, 245
- Island Hopping, 258
- Island Model, 258
- Isolated, 47
- Isolation
 - by distance, 262
- Iteration, 34
- IWLCS, 198
- Java, 391
- java.io.Serializable, 394
- java.util.Random, 461, 463
- JavaGroups, 364
- JB, 139
- Kauffman NK, 281
- Kernel Density Estimation, 458
- KES, 62
- Kleene closure, 502
- Kleene star, 502
- Kurtosis, 427
 - excess, 427
- LAN, 491
- Landscape
 - fitness, 29
 - problem, 29
- Language, 501, 502
 - formal, 502
 - guarded command, 172
- Language Attribute, 509
- Laplace
 - assumption, 418
- Large Numbers
 - law of, 429
- Las Vegas Algorithm, 484
- Latency, 499
- LCG, 460
- LCS, 72, 171, 172, 197, 198, 322
 - Michigan-style, 207
 - Pitt, 206
 - Pittsburgh-style, 207
- LDSE, 226
- Learning
 - machine, 128, 142, 195, 215, 219, 230, 239
- Learning Classifier System
 - Michigan-style, 207
 - Pittsburgh-style, 207
- Learning Classifier Systems, 72, 197, 198, 322
- LEFT, 305
- Left-total, 413
- Length
 - defined, 122
- Levels-back, 166
- Lexeme, 157, 502
- LGP, 159, 160, 363
 - page-based, 161
- LGPL, 533
- License, VI, 527, 533
 - FDL, 527

- LGPL, 533
- Life
 - artificial, 178
- Lifting, 135
- Likelihood, 451
 - function, 451
- Linear Congruential Generator, 460
- Linear Genetic Programming
 - page-based, 161
- Linear Order, 415
- Linkage
 - Average, 471
 - Complete, 471
 - Single, 471
- List, 411
 - addListItem, 411
 - appendList, 411
 - countItemOccurrences, 411
 - createList, 411
 - deleteListItem, 411
 - deleteListRange, 411
 - insertListItem, 411
 - listToSet, 413
 - removeListItem, 413
 - setToList, 413
 - sortList, 412
 - subList, 412
 - search (sorted), 412
 - search (unsorted), 412
 - sorting, 412
- LLN, 429
- Local Search, 246
- LocalAnt, 357
- Locality, 40, 59
- Locus, 24
- LOGENPRO, 147
- LS-1, 206
- LTAG, 511
- Lunch
 - no free, 54
- MA, 241–243
- Machine
 - finite state, 126, 128, 195, 304, 359, 361
 - learning, 128, 142, 195, 215, 219, 230, 239
- Manhattan Distance, 469
- Mapping
 - Genotype-Phenotype, 121, 138
 - genotype-phenotype, 25
- Mask, 122
 - defined length, 122
 - order, 122
- Master-Slave, 257
- matchesCondition, 200
- Mathematics, 61, 77, 152, 159
- Maximum, 7, 370, 423
 - global, 8
 - local, 7
- Maximum Likelihood Estimator, 453
- MCDM, 61, 63
- Mean, 368
 - arithmetic, 424
- Median, 427, 428
- Medicine, 77, 114, 128, 197, 215, 467
- Memetic Algorithm, 242, 243
- Memetic Algorithms, 241
- Memory
 - with memory, 178
- Memory Consumption, 246
- Mendel, 63, 80, 115, 129
- mergeAction, 202
- Message, 198
- Metaheuristic, 5
- Method of Inequalities, 16
- Metropolis, 217
- MIC, 63
- MICA2, 492
- MICAI, 64
- Microcontroller, 491
- migrate, 258, 259
- Military, 77
- Minimization, 7
- Minimum, 7, 370, 423
 - global, 8
 - local, 7
- MLE, 453
- Model, 36
- Module Mutation, 167
- MOEA, 68
- Moment, 426
 - about mean, 426
 - central, 426
 - standardized, 426
 - statistical, 426
- Monotone
 - function, 414
 - heuristic, 252
- Monotonic, 414
- Monotonicity, 414
- Monte Carlo
 - method, 484
- Monte Carlo Algorithm, 484
- MOVE, 305
- MPX, 119
- MSB, 492
- MSE, 450
- Multi-objective, 68, 186, 189, 392

- multi-objective, 9
- Multi-Objectivity, 294
- Multicast, 497
- Multimodality, 37
- Mutation, 70, 108, 118–120, 131, 132, 164
 - global, 164
 - link, 164
 - module, 167
 - path, 354
 - tree, 131
- Natural Representation, 117
- NCGA, 109
- ND, 43–45, 285
- ND Landscape, 285
- Needle-In-A-Haystack, 47, 167, 186
- Negative Slope Coefficient, 42
- Network
 - butterfly, 372
 - European optical, 349
 - neutral, 46
 - of sensors, 491
 - on chip, 359
 - random Boolean, 46
 - self-organized, 349
 - selfish, 349
 - sensor, 365
 - wireless, 365
 - singular, 349
 - wireless, 365
- Network Topology, 486
- Networking, 61, 77, 114, 128, 210, 219, 230, 234, 242, 247
- Neural Network
 - artificial, 318
- Neutral
 - network, 46
- Neutrality, 43, 166, 167, 187, 283, 285, 286, 293, 301
 - explicit, 166
 - implicit, 167
- nextGaussian, 463
- NFL, 54
- NIAH, 47
- Niche Count, 86
- Niche Preemption Principle, 37
- NK, 40, 281, 284
- NKp, 45, 283, 284
- NKq, 45, 283, 284
- No Free Lunch, 54
- NoC, 359
- Node Selection, 136
- nodeWeight, 137
- Noise, 49, 50, 53
- Non-Decreasing, 414
- Non-functional, 184, 186
- Non-Increasing, 415
- Non-Synonymous Change, 45
- Nonile, 428
- Norm, 470
 - Euclidian, 470
 - infinity, 470
 - Manhattan, 470
 - p, 470
- Normal Distribution, 437
 - standard, 437
- NSC, 42
- Nucleus, 471
- Numbers
 - complex, 408
 - integer, 408
 - law of large, 429
 - natural, 408
 - pseudorandom, 460
 - random, 459
 - rational, 408
 - real, 408
- NWGA, 115
- OAB, 359
- OAS, 359
- Objective Function, 3
- Objective Space, 27
- ObjectivePrinterPipe, 403
- One-Fifth Rule, 193
- OneMax, 289
- ontogenic mapping, 25
- Operations Research, 61, 230, 239
- Operator
 - path genetic, 354, 355
- Operator Correlation, 42
- Optics, 61, 191
- Optimality
 - Pareto, 13
- Optimiality, 246
- Optimization, 30
 - multi-objective, 9
 - algorithm, 30
 - global, 30
 - global, 3
 - algorithm, 30
 - taxonomy, 4
 - multi-objective, 9, 189, 392
 - Pareto, 13
 - prevalence, 20
 - weighted sum, 11
- offline, 6
- online, 6

- random, 213
 - termination criterion, 35
- Optimum, 7, 27
 - global, 8
 - isolated, 47
 - local, 7, 27
 - optimal set, 8, 10
 - extracting, 264
 - obtain by deletion, 264
 - pruning, 265
 - updating, 263
 - updating by insertion, 263
- Order, 122
 - linear, 415
 - partial, 13, 415
 - simple, 415
 - total, 415
- org.sigoa, 394
- org.sigoa.refimpl, 394
- org.sigoa.spec, 394
- OSPF, 354
- Output, 481
- Overfitting, 48, 50, 53, 188, 292, 343
- Overgeneralization, 51
- Overlay Network, 486
- Oversimplification, 51, 292
- OWL-S, 328
- p-Spin, 40, 284
- P2P, 259, 353, 489
- PADO, 136, 161, 162
- PAES, 266
- Parallel Algorithm Discovery and Orchestration, 161
- Parallelism
 - implicit, 70
 - intrinsic, 70
- Parallelization, 255
- Pareto, 13
 - frontier, 13
 - optimal, 13
 - ranking, 85
 - set, 13
- Partial Order, 415
- Partial order, 13
- Particle Swarm Optimization, 237
- Partition
 - static, 497
- Parzen window, 458
- Path
 - Crossover, 354
 - genetic operator, 354, 355
 - mutation, 354
- Pattern Recognition, 128, 467
- PDF, 422
- PDGP, 162
- Peer-To-Peer, 259
- Peer-to-Peer, 353, 489
- Percentile, 428
- Permutation, 119, 133, 419
 - tree, 132
- Perturbation, 52
- Petri Net
 - with registers, 359
- Phenome, 23
- Phenotype, 23
- Phosphate, 24
- Physics, 114, 128, 142, 191, 219, 224, 239
- Pitt approach, 206, 322
- PL, 139
- PMF, 422
- PNR, 359
- Point
 - crossover, 119
- Point Estimator, 450
- Poisson, 431
 - Process, 432
- Poisson Distribution, 431
- Population, 28, 67
- population, 73
- PPSN, 80
- Prediction
 - sequence, 197
- Preemption
 - niche, 37
- Prehistory, 431
- Premature Convergence, 37
- Preservative Selection, 74
- Prevalence, 20
- Prevention
 - convergence, 107
- Principle
 - niche preemption, 37
- Probabilistic Algorithm, 484
- Probability
 - von Mises [1717], 419
 - Bernoulli, 418
 - Conditional, 420
 - Kolmogorov, 419, 420
 - of success, 42, 44
 - space, 420
 - of a random variable, 421
- Probability Density Function, 422
- Probability Mass Function, 422
- Probit, 439
- Problem
 - Terminal Assignment, 348

- Problem Landscape, 29
- Problem Space, 23
- Processing, 481
- Product
 - Cartesian, 410
- Production Systems, 197
- PROGN2, 305
- PROGN3, 305
- Programming
 - automated, 128, 152
- PROLOG, 147
- Property
 - constant innovation, 46
- Protocol, 359
- Protocols
 - Aggregation, 366, 367
 - gossip-based, 367
 - proactive, 367
 - reactive, 367
- Pruning, 265
- Pseudorandom Numbers, 460
- PSO, 226, 237, 238
 - hybrid simplex, 226
- psoReproduce, 238
- Push, 179
- Push3, 179
- PushGP, 179, 180
- Pushpop, 179, 180

- QBB, 207
- QoS, 339
- Quality of Service, 339
- Quantile, 428
- Quartile, 428
- Quine, 182
- Quintile, 428

- Radio, 491
- RAG, 153, 509
- rag, 153
- Ramped Half-and-Half, 130
- Random
 - event, 417
 - Experiment, 417
 - experiment, 420
 - variable, 421
 - continuous, 422
 - discrete, 422
- Random Boolean Network, 46
- random neighbors, 282
- Random Number
 - generator
 - normally distributed, 461
 - uniformly distributed, 461
- Random Numbers, 459
 - pseudo, 460
 - uniformly distributed, 460
- Random Optimization, 213, 214
 - heuristic, 214
- Random Walk, 42, 43, 250
- Randomized Algorithm, 484
- Range, 423
 - interquartile, 428
- Ranking
 - Pareto, 85
- RBGP, 171–175, 308, 312
- RBN, 46
- Real-Encoded, 117
- Recognizers, 502
- Recombination, 70, 108
 - tree, 131
- Recursive Adaptive Grammars, 509
- Redundancy, 26, 46, 293
- Reflexivenss, 415
- Register
 - connection, 170
- Regression, 341
 - logistic, 318
 - Symbolic, 341
 - symbolic, 128, 142, 159, 167
- Relation
 - binary, 413
 - types and properties of, 413
 - equivalence, 416
 - order, 415
 - partial, 415
 - total, 415
- Reliability, 498
 - message delay, 499
 - message latency, 499
 - message loss, 498
 - message modification, 499
- Repair, 145
- Representation
 - natural, 117
 - state-based, 361
- Reproduction, 108
 - asexual, 117
 - NCGA, 110
 - path genetic, 354, 355
 - sexual, 67, 70, 117
- Ressources, 77, 191
- REVOP, 73
- RFS, 350
- RIGHT, 305
- Ring, 488
- RIP, 354

- Road
 - royal, 285
 - real, 288
 - variable-length, 287
 - VLR, 287
- Robotics, 77, 128, 152, 167, 195, 210, 224
- Robustness, 51, 53
- Routing, 352
 - genetic adaptive, 354
 - source, 354
- Royal Road, 45, 285, 287
 - real, 288
 - variable-length, 287
 - VLR, 287
- Royal Tree, 288
- Ruggedness, 40, 284, 294, 298, 302
- Rule
 - semantic, 505
- Rule-based Genetic Programming, 171
 - extended, 176
- S-Expressions, 512
- SA, 217, 219
- SAAN, 163
- Salience
 - high, 289
 - low, 289
- Sample space, 417
- Santa Fe trail, 303
- SBR, 361
- Scalability, 486
- ScatterNode, 492
- Scatterweb, 492
- Scheduling, 77, 114, 191, 194, 234, 242
- Schema, 42, 56, 122
 - Theorem, 122
 - theorem, 123
- Schema Theorem, 43, 285
- Science
 - computer
 - theoretical, 479
- Search
 - A*, 252
 - best-first, 251
 - breadth-first, 247
 - depth-first, 248
 - iterative deepening, 249, 363
 - depth-limited, 249
 - greedy, 251
 - informed, 251
 - local, 246
 - State Space, 245
 - uninformed, 247
- Search Operations, 24
- Search Space, 23
- Selection, 92
 - Deterministic, 94
 - elitist, 75
 - extinctive, 73
 - left, 74
 - right, 74
 - extrema, 46
 - Fitness Proportionate, 95, 97, 98
 - Node, 136
 - Ordered, 103
 - ordered
 - with replacement, 104
 - without replacement, 104
 - preservative, 74
 - ranking, 105
 - Roulette Wheel, 95, 97, 98
 - Threshold, 94
 - Tournament, 100
 - non-deterministic, 103
 - with replacement, 101, 103
 - without replacement, 102
 - Truncation, 94
 - truncation, 94
 - vega, 105, 106
- Self-Modification, 179
- Self-modification, 362
- Self-Organization, 349
- Selfishness, 349
- Semantic, 501
- Semantic Rule, 505
- Sensor Network, 365, 491
 - wireless, 491
- Sentence, 502
- Server, 488
- Service Oriented Architecture, 327
- Set, 407
 - cardinality, 407
 - Cartesian product, 410
 - complement, 410
 - countable, 410
 - difference, 409
 - empty, 408
 - equality, 407
 - intersection, 409
 - List, 411
 - membership, 407
 - operations on, 408
 - optimal, 8, 10
 - Power set, 410
 - relations between, 407
 - special, 408
 - subset, 408

- superset, 408
- theory, 407
- Tuple, 410
- uncountable, 410
- union, 409
- home, 490
- sexp, 512
- Sexual Reproduction, 67, 70, 117
- SGP, 126
- SH, 212
- Sharing
 - Variety Preserving, 87
- Sharing Function, 85
- Sigoa, 391, 397
- Simple Order, 415
- Simplex
 - concurrent, 227
 - downhill, 223
- Simplex Evolution
 - low dimensional, 226
- Simulated Annealing, 217, 219
 - simulated quenching, 219
 - temperature schedule, 219
- Simulated Quenching, 219
- Simulation, 36
- Simulation Dynamics, 42
- SIS, 239
- Skewness, 427
- Small- $\mathbf{o}$ notation, 483
- Small- ω notation, 484
- SMART, 162
- Smart Home, 365
- SmartMesh, 492
- SOA, 327, 328
- Soft Assignment, 178
- Software engineering, 185
- Software testing, 185
- Solution Candidate, 23
- Solution Space, 23
- Space
 - objective, 27
 - problem, 23
 - search, 23
 - solution, 23
- Spin-Glass, 29, 284
- SPX, 119
- SSAAN, 164
- SSEA, 74
- SSIAN, 164
- Standard Deviation, 425
- Standard Genetic Programming, 126
- Star, 488
- State Machine
 - finite, 126, 195, 304
- State Space
 - Search, 245
- Statistical Independence, 420
- Statistics, 423
- Steady State, 74
- STGP, 146, 147
- Sticky Crossover, 161
- Stigmergy, 233
 - sematectonic, 233
 - sign-based, 233
- Stochastic
 - Theory, 417
- Stochastic Gradient Descent, 212
- Stopping Criterion, 35
- String, 502
 - character, 502
- Strongly Typed Genetic Programming, 146
- Student's t-Distribution, 445
- Substitution, 510
- Success
 - probability of, 42, 44, 277
- Sugar, 24
- Sum, 370, 424
 - sqr, 425
- Support Vector Machine, 318
- Surjective, 413
- Surjectivity, 413
- SVM, 318
- Symbol
 - grammar, 503
 - non-terminal, 503
 - start, 503
 - terminal, 503
- Symbolic Regression, 128, 142, 159, 167, 341
 - Example, 343
- Symmetric, 416
- Symmetry, 416
- Synonymous Change, 45
- Syntax, 501
- t-Distribution, 445
- TA, 348
- Tabu Search, 229
 - multi-objective, 230
- TAG, 156, 510
 - lexicalized, 511
 - LTAG, 511
- TAG3P, 156–159
- talk, 490
- TB, 140
- Technological Landscape, 284
- Terminal Assignment Problem, 348
- Termination, 482

- Termination Criterion, 35
- Ternary System, 200
- TGP, 126, 127, 130, 160
- Theorem
 - Schema, 122, 285
- Threshold
 - error, 42, 290
- Threshold Selection, 94
- Thymine, 24
- Time Consumption, 246
- TODO, VI, 7, 18, 53, 92, 107, 111, 280, 457, 459, 486, 515
- Toplogy, 348
- Topology, 486, 496
 - bus, 488
 - fully connected, 488
 - grid, 488
 - hierarchical, 488
 - hierarchy, 488
 - network, 486
 - ring, 488
 - star, 488
 - unrestricted, 487
- Total Order, 415
- TPX, 119
- Transitivity, 413
- Transsititve, 413
- Trap Function, 43, 285
- Tree
 - auxiliary, 510, 511
 - decision, 318
 - derivation, 503
 - elementary, 511
 - initial, 511
 - royal, 288
- Tree Genomes, 130
- Tree-Adjoining Grammar, 510
- Truncation Selection, 94
- TS, 229
- TSoR, 100
- TSR, 100
- Tuple, 410
 - Type, 410
- Type, 410
- Type-0, 503
- Type-1, 503
- Type-2, 503
- Type-3, 503
- Types Possibilities Tables, 147
- Ubiquitous Computing, 365
- Unicast, 497
- Uniform Distribution
 - continuous, 436
 - discrete, 430
- uniformSelectNode, 137, 138
- Uninformed Search, 247
- Variable, 503
- Variance, 370, 424
 - epistatic, 42
 - estimator, 425
- Variety Preserving, 87
- VEGA, 110
- Walk
 - Adaptive, 252
 - adaptive
 - fitter dynamics, 253
 - greedy dynamics, 253
 - one-mutant, 252
 - drunkyard's, 250
 - random, 42, 43, 250
- Wasp, 199
- Web Browser, 489
- Web Server, 489
- Web Service, 327
 - Challenge, 327
- Web Service Challenge, 328
- Website, 489
- Weighted Sum, 11, 20, 83
- Wildcard, 122, 200, 322, 326
- Wireless LAN, 491
- Wireless Sensor Network, 491
- WLAN, 363
- Wrapping, 134
- WS-Challenge, 328
- WSC, 328, 339
- WSDL, 328
- WWW, 489
- XCS, 207
- ZCS, 207

List of Figures

1.1	The taxonomy of global optimization algorithms.	5
1.2	Global and local optima of a two-dimensional function.	8
1.3	Two functions f_1 and f_2 with different maxima $\hat{\mathbf{x}}_1$ and $\hat{\mathbf{x}}_2$	10
1.4	Two functions f_3 and f_4 with different minima $\check{\mathbf{x}}_1$, $\check{\mathbf{x}}_2$, $\check{\mathbf{x}}_3$, and $\check{\mathbf{x}}_4$	10
1.5	Optimization using the weighted sum approach (first example).	11
1.6	Optimization using the weighted sum approach (second example).	12
1.7	A problematic constellation for the weighted sum approach.	12
1.8	Optimization using the Pareto Frontier approach.	14
1.9	Optimization using the Pareto Frontier approach (second example).	15
1.10	Optimization using the Method of Inequalities approach (first example).	17
1.11	Optimization using the Method of Inequalities approach (first example).	17
1.12	An external decision maker providing an evolutionary algorithm with utility values.	19
1.13	Spaces, Sets, and Elements involved in an optimization process.	22
1.14	A sketch of a part of a DNA molecule.	24
1.15	The relation of genome, genes, and the problem space.	26
1.16	An example optimization problem.	31
1.17	The problem landscape of the example problem derived with searchOp ₁	32
1.18	The problem landscape of the example problem derived with searchOp ₂	33
1.19	Premature convergence in the objective space.	38
1.20	Different possible properties of fitness landscapes.	41
1.21	Possible positive influence of neutrality.	45
1.22	The influence of epistasis on the fitness landscape.	48
1.23	Overfitting due to complexity.	49
1.24	Fitting noise.	50
1.25	Oversimplification.	51
1.26	A robust local optimum vs. a “unstable” global optimum.	52
1.27	A visualization of the No Free Lunch Theorem.	55
1.28	An graph coloring-based example for properties and formae.	56
1.29	Example for formae in symbolic regression.	57
2.1	The basic cycle of evolutionary algorithms.	68
2.2	The family of evolutionary algorithms.	73
2.3	The configuration parameters of evolutionary algorithms.	76
2.4	An example scenario for Pareto ranking.	84
2.5	The sharing potential in the Variety Preserving Ranking example.	90
2.6	The two example fitness cases.	94
2.7	The number of expected offspring in truncation selection.	95
2.8	Examples for the idea of roulette wheel selection.	96

2.9	The number of expected offspring in roulette wheel selection.	99
2.10	The number of expected offspring in tournament selection.	101
2.11	The number of expected offspring in ordered selection.	105
2.12	The expected numbers of occurrences for different values of n and cp	107
3.1	The basic cycle of genetic algorithms.	114
3.2	A four bit string genome $\mathbb{G}$ and a fictitious phenotype $\mathbb{X}$	119
3.3	Value-altering mutation of string chromosomes.	120
3.4	Permutation applied to a string chromosome.	120
3.5	Crossover (recombination) of fixed-length string chromosomes.	121
3.6	Mutation of variable-length string chromosomes.	121
3.7	Crossover of variable-length string chromosomes.	122
3.8	An example for schemata in a three bit genome.	124
4.1	Genetic Programming in the context of the IPO model.	127
4.2	The AST representation of algorithms/programs.	129
4.3	Tree creation by the <i>full</i> method.	133
4.4	Tree creation by the <i>grow</i> method.	133
4.5	Possible tree mutation operations.	134
4.6	Tree crossover by exchanging sub-trees.	134
4.7	Tree permutation – asexually shuffling sub-trees.	135
4.8	Tree editing – asexual optimization.	135
4.9	An example for tree encapsulation.	136
4.10	An example for tree wrapping.	136
4.11	An example for tree lifting.	137
4.12	Automatically defined functions in Genetic Programming.	138
4.13	Comparison of functions and macros.	139
4.14	A GPM example for Gene Expression Programming.	143
4.15	An example for edge encoding.	146
4.16	Example for valid and invalid trees in symbolic regression.	147
4.17	Example for valid and invalid trees in typed Genetic Programming.	148
4.18	The structure of a Grammatical Evolution system [1000].	152
4.19	An TAG realization of the C-grammar of Listing 4.6.	158
4.20	One example genotype-phenotype mapping in TAG3P.	160
4.21	The impact of insertion operations in Genetic Programming	162
4.22	The general structure of a Parallel Algorithm Discovery and Orchestration program.	164
4.23	The term $\max\{x * y, x * y + 3\}$	164
4.24	An example program in Genetic Network Programming syntax.	167
4.25	An example for the GPM in Cartesian Genetic Programming.	168
4.26	Epistasis in a Grammatical Evolution-like approach.	170
4.27	Positional epistasis in Genetic Programming.	172
4.28	The difference between linear Genetic Programming and Algorithmic Chemistries.	173
4.29	genotype-phenotype mapping in Rule-based Genetic Programming.	174
4.30	A program computing the faculty of a natural number in Java and RBGP syntax.	176
4.31	The tree phenotype (and genotype) of Listing 4.15.	179
4.32	A Fraglet-based election algorithm.	185
4.33	A simple quine Fraglet (borrowed from [1095])	185
7.1	The structure of a Michigan style Learning Classifier System according to Geyer-Schulz [1225].	199
7.2	One possible encoding of messages for a frog classifier system	200

11.1 One possible step of the downhill simplex algorithm applied to a problem in $\mathbb{R}^2$.	226
17.1 Parallelization potential in evolutionary algorithm.	256
17.2 A sequentially proceeding evolutionary algorithm.	256
17.3 A parallel evolutionary algorithm with two worker threads.	257
17.4 An EA distributed according to the client-server approach.	258
17.5 An evolutionary algorithm distributed in a P2P network.	259
17.6 An example for a heterogeneous search.	260
17.7 A mixed distributed evolutionary algorithms.	261
20.1 An example for the moving peaks benchmark of Branke [227]	281
20.2 The perfect Royal Trees.	288
20.3 Example fitness evaluation of Royal Trees	289
20.4 An example for the fitness landscape model.	291
20.5 An example for the epistasis mapping $z \rightarrow e_4(z)$.	293
20.6 An example for r_γ with $\gamma = 0..10$ and $\hat{f} = 5$.	296
20.7 The basic problem hardness.	297
20.8 Experimental results for the ruggedness.	298
20.9 Experiments with ruggedness and deceptiveness.	300
20.10 Experiments with epistasis.	300
20.11 The results from experiments with neutrality.	301
20.12 Expectation and reality: Experiments involving both, epistasis and neutrality	302
20.13 Expectation and reality: Experiments involving both, ruggedness and epistasis	302
20.14 The Santa Fee Trail in the Artificial Ant Problem [65].	304
20.15 The f_1 /generation-plots of the best configurations.	313
21.1 Some logos of the DATA-MINING-CUP.	318
21.2 A few samples from the DMC 2007 training data.	320
21.3 DMC 2007 sample data – same features but different classes.	321
21.4 An example classifier for the 2007 DMC.	323
21.5 The course of the classifier system evolution.	324
21.6 Some Pareto-optimal individuals among the evolved classifier systems.	325
21.7 The course of the modified classifier system evolution.	326
21.8 The logo of the Web Service Challenge.	329
21.9 A sketch of the Pareto front in the genetic composition algorithm.	335
21.10 The WSC 2007 Composition System of Bleul et al. [1604; 1605].	337
21.11 The Knowledge Base and Service Registry of our Composition System.	338
22.1 An example genotype of symbolic regression of with $x = \mathbf{x} \in \mathbb{R}^1$.	342
22.2 $\varphi(x)$, the evolved $\psi_1^*(x) \equiv \varphi(x)$, and $\psi_2^*(x)$.	345
23.1 The initial network design of the EON network [712, 1629].	350
23.2 The genotype representing the connection matrix of the nine central nodes of EON.	350
23.3 Snippets from the GP representations of the bit string from Figure 23.2	351
23.4 An evolved routing rule set (left) with target node 0 for the specified network (right).	354
23.5 The table with the most frequently used routes of node 0.	354
23.6 The path genetic operations in GARA.	355
23.7 The application of Listing 23.1 to a 5×5 grid.	358
23.8 The scenario for protocol synthesis as used in [859].	360
23.9 The best individual from the experiment of de Araújo et al. [859].	361
23.10 Global $\rightarrow$ Local behavior transformations.	366

23.11	The two basic forms of aggregation protocols.	368
23.12	An example sensor network measuring the temperature.	369
23.13	An gossip-based aggregation of the average example.	369
23.14	Optimal data dissemination strategies.	372
23.15	The model of a node capable to execute a proactive aggregation protocol. ...	373
23.16	The behavior of the distributed average protocol in different scenarios.	378
23.17	A dynamic aggregation protocol for the distributed average.	379
23.18	Some examples for the formula series part of aggregation protocols.	380
23.19	The evolutionary progress of the static <i>average</i> protocol.	383
23.20	The relation of f_1 and f_2 in the static <i>average</i> protocol.	384
23.21	The evolutionary progress and one grown solution of the static <i>root-of-average</i> protocol.	384
23.22	The relation of f_1 and f_2 in the static <i>root-of-average</i> protocol.	385
23.23	The evolutionary progress of the dynamic <i>average</i> protocol.	386
23.24	The relation of f_1 and f_2 in the dynamic <i>average</i> protocol.	386
23.25	The evolutionary progress and one grown solution of the dynamic <i>root-of-average</i> protocol.	387
23.26	The relation of f_1 and f_2 in the dynamic <i>root-of-average</i> protocol.	388
24.1	The top-level packages of the Sigoa optimization system.	394
24.2	The subsystem specification of the optimization framework.	395
26.1	Set operations performed on sets A and B inside a set $\mathbb{A}$	409
26.2	Properties of a binary relation R with domain A and codomain B	414
27.1	The PMFs of some discrete uniform distributions.	431
27.2	The CDFs of some discrete uniform distributions.	431
27.3	The PMFs of some Poisson distributions.	432
27.4	The CDFs of some Poisson distributions.	433
27.5	The PMFs of some binomial distributions.	435
27.6	The CDFs of some binomial distributions.	435
27.7	The PDFs of some continuous uniform distributions.	436
27.8	The CDFs of some continuous uniform distributions.	437
27.9	The PDFs of some normal distributions.	438
27.10	The CDFs of some normal distributions.	438
27.11	The PDFs of some exponential distributions.	441
27.12	The CDFs of some exponential distributions.	441
27.13	The PDFs of some χ^2 distributions.	442
27.14	The CDFs of some χ^2 distributions.	443
27.15	The PDFs of some Student's t-distributions.	446
27.16	The CDFs of some Student's t-distributions.	446
27.17	The PMF and CMF of the dice throw.	448
27.18	The numbers thrown in the dice example.	449
28.1	A clustering algorithm applied to a two-dimensional dataset A	467
29.1	The relation between algorithms and programs.	480
29.2	A process in the IPO model.	481
29.3	Some simple network topologies.	487
29.4	Multiple clients connected with one server.	489
29.5	A peer-to-peer system in an unstructured network.	490
29.6	A block diagram outlining building blocks of a sensor node.	491
29.7	Images of some sensor network platforms.	493
29.8	Synchronous parallelism in a model of a network of five nodes.	495

29.9 Asynchronous parallelism in a model of a network of five nodes.	496
29.10 Dynamic topology due to overlapping active times of nodes.	497
29.11 The three different message transmission types.	498
29.12 The derivation of the example expansion of the grammar G	504
29.13 An instantiation of the grammar from Listing 29.7.	506
29.14 One possible expansion of the example grammar G_2	508
29.15 An example TAG tree.	510
29.16 An example for the substitution operation.	511
29.17 An example for the adjunction operation.	512

List of Tables

2.1	The Pareto domination relation of the individuals illustrated in Figure 2.4. . .	84
2.2	An example for Variety Preserving Ranking based on Figure 2.4.	90
2.3	The distance and sharing matrix of the example from Table 2.2.	91
4.1	One possible operator set of edge encoding.	145
4.2	Some Fraglet instructions (from [1092] and http://www.fraglets.net/). ...	184
7.1	if-then rules for frogs	199
7.2	if-then rules for frogs in encoded form	202
19.1	The basic optimization problem settings.	272
19.2	Some basic optimization algorithm settings.	273
19.3	Some additional parameters of experiments.	274
19.4	Some basic measures that can be obtained from experiments.	275
19.5	Simple evaluation results.	277
20.1	The Sphere function.	280
20.2	The settings of the experiments with the benchmark model.	296
20.3	First-level evaluation results of the experiments with the model benchmark. .	297
20.4	The settings of the RBGP-Genetic Programming experiments for the GCD problem.	309
20.5	Evaluation parameters used in Table 20.6.	310
20.6	Results of the RBGP test series on the GCD problem.	311
21.1	Feature-values in the 2007 DMC training sets.	321
21.2	Feature conditions in the rules.	322
21.3	The settings of the experiments for the DATA-MINING-CUP.	324
21.4	Different feature conditions in the rules.	326
21.5	Experimental results for the web service composers.	336
22.1	Sample Data $A = \{(x_i, y_i) : i \in [0..8]\}$ for Equation 22.8	344
23.1	The Node-Types of the NP-approaches.	352
23.2	The settings of the Aggregation-Genetic Programming experiments.	382
27.1	Special Quantiles	428
27.2	Parameters of the discrete uniform distribution.	430
27.3	Parameters of the Poisson distribution.	432
27.4	Parameters of the Binomial distribution.	434
27.5	Parameters of the continuous uniform distribution.	436

738 LIST OF TABLES

27.6	Parameters of the normal distribution.	437
27.7	Some values of the standardized normal distribution.....	439
27.8	Parameters of the exponential distribution.	440
27.9	Parameters of the χ^2 distribution.	442
27.10	Some values of the χ^2 distribution.	444
27.11	Parameters of the Student's t- distribution.	445
27.12	Table of Student's t-distribution with right-tail probabilities.....	447
27.13	Parameters of the dice throw experiment.....	449
29.1	Some examples of the big- O notation	483
29.2	The Chomsky Hierarchy.....	503

List of Algorithms

1.1	Example Iterative Algorithm	35
2.1	$X^* \leftarrow \text{simpleEA}(\text{cmp}_F, ps)$	72
2.2	$X^* \leftarrow \text{elitistEA}(\text{cmp}_F, ps, a)$	75
2.3	$v \leftarrow \text{assignFitnessParetoRank}(Pop, \text{cmp}_F)$	85
2.4	$v \leftarrow \text{assignFitnessVarietyPreserving}(Pop, \text{cmp}_F)$	88
2.5	$v \leftarrow \text{assignFitnessTournament}_{q,r}(Pop, \text{cmp}_F)$	92
2.6	$Mate \leftarrow \text{truncationSelect}_k(Pop, v, ms)$	94
2.7	$Mate \leftarrow \text{rouletteWheelSelect}_r(Pop, v, ms)$	97
2.8	$Mate \leftarrow \text{rouletteWheelSelect}_w(Pop, v, ms)$	98
2.9	$Mate \leftarrow \text{tournamentSelect}_{r,k}(Pop, v, ms)$	101
2.10	$Mate \leftarrow \text{tournamentSelect}_{w1,k}(Pop, v, ms)$	102
2.11	$Mate \leftarrow \text{tournamentSelect}_{w2,k}(Pop, v, ms)$	102
2.12	$Mate \leftarrow \text{tournamentSelect}_{r,k}^P(Pop, v, ms)$	103
2.13	$Mate \leftarrow \text{orderedSelect}_r^P(Pop, v, ms)$	104
2.14	$Mate \leftarrow \text{orderedSelect}_w^P(Pop, v, ms)$	104
2.15	$Mate \leftarrow \text{vegaSelectionPopFms}$	106
2.16	$Pop' \leftarrow \text{convergencePrevention}(Pop, cp)$	107
2.17	$Pop \leftarrow \text{createPop}(s)$	109
2.18	$Pop \leftarrow \text{nsgaReproducePop}_{foc}(Mate)$	110
2.19	$X^* \leftarrow \text{vega}(F, s)$	111
4.1	$n \leftarrow \text{uniformSelectNode}(t)$	140
4.2	<i>Halting Problem: reductio ad absurdum</i>	187
7.1	$S \leftarrow \text{matchesConditions}(M, C)$	201
7.2	$\text{nonLearningClassifierSystem}(P)$	203
7.3	$\text{learningClassifierSystem}()$	205
8.1	$x^* \leftarrow \text{hillClimber}(f)$	210
8.2	$x^* \leftarrow \text{hillClimberMO}(\text{cmp}_F, a)$	211
8.3	$X^* \leftarrow \text{hillClimberMO\_RR}(\text{cmp}_F, a)$	212
9.1	$x^* \leftarrow \text{randomOptimizer}f$	214
10.1	$x^* \leftarrow \text{simulatedAnnealing}(f)$	218
10.2	$x^* \leftarrow \text{simulatedAnnealingMO}(\text{cmp}_F, a)$	221
11.1	$x^* \leftarrow \text{downhillSimplex}(f)$	225
12.1	$x^* \leftarrow \text{tabuSearch}(f, n)$	230
12.2	$x^* \leftarrow \text{tabuSearchMO}(\text{cmp}_F, n, a)$	231
14.1	$x^* \leftarrow \text{psoOptimizer}fps$	238
15.1	$Pop \leftarrow \text{createPopMA}(s, \text{cmp}_F)$	242
15.2	$Pop \leftarrow \text{reproducePopMA}(Mate, \text{cmp}_F)$	242
16.1	$P \leftarrow \text{expandToInd}(g)$	246
16.2	$X^* \leftarrow \text{bfs}(r, \text{isGoal})$	248

16.3	$X^* \leftarrow \text{dfs}(r, \text{isGoal})$	249
16.4	$X^* \leftarrow \text{dl\_dfs}(r, \text{isGoal}, d)$	249
16.5	$X^* \leftarrow \text{iddfs}(r, \text{isGoal})$	250
16.6	$X^* \leftarrow \text{greedySearch}(r, \text{isGoal}, h)$	252
18.1	$P_{new}^* \leftarrow \text{updateOptimalSet}(P_{old}^*, p_{new})$	264
18.2	$P_{new}^* \leftarrow \text{updateOptimalSet}(P_{old}^*, p_{new})$ (<i>2nd Version</i>)	264
18.3	$P_{new}^* \leftarrow \text{updateOptimalSetN}(P_{old}^*, P)$	264
18.4	$P^* \leftarrow \text{extractOptimalSet}(Pop)$	265
18.5	$P_{new}^* \leftarrow \text{pruneOptimalSet}_c(P_{old}^*, k)$	266
18.6	$(P_l, lst, cnt) \leftarrow \text{agaDivide}(P_{old}, d)$	267
18.7	$P_{new}^* \leftarrow \text{pruneOptimalSet}_{aga}(P_{old}^*, d, k)$	268
20.1	$r_\gamma \leftarrow \text{buildRPermutation}(\gamma, \hat{f})$	295
20.2	$\gamma \leftarrow \text{translate}(\gamma', \hat{f})$	299
20.3	$\text{gcd}(a, b) \leftarrow \text{euclidGcdOrig}(a, b)$	306
20.4	$\text{gcd}(a, b) \leftarrow \text{euclidGcd}(a, b)$	306
20.5	$r \leftarrow f_1(x, a, b)$	307
21.1	$S \leftarrow \text{webServiceCompositionIDDFS}(R)$	332
21.2	$r \leftarrow \text{cmp}_{wsc}(S_1, S_2)$	333
21.3	$S \leftarrow \text{greedyComposition}(R)$	334
23.1	$\text{gossipBasedAggregation}()$	368
23.2	$\text{simulateNetwork}(m, T)$	375
23.3	$f_1(u, e, r) \leftarrow \text{evaluateAggregationProtocol}(u, m, T)$	376
27.1	$\text{dist}_{cr}^\rho(\dots, A) \leftarrow \text{computeCrowdingDistance}(a, A)$	459
27.2	$r \leftarrow \text{random}_l(\text{random}, \text{low}, \text{high})$	462
27.3	$(n_1, n_2) \leftarrow \text{random}_{n, p^2}()$	463
27.4	$y \leftarrow \text{random}_{bs}(\mu, \sigma)$	464
28.1	$C_{new} \leftarrow \text{kMeansModify}_k(C)$	473
28.2	$C \leftarrow \text{kMeansCluster}_k(A)$	474
28.3	$C \leftarrow \text{nNearestNeighborCluster}_k(n) A$	475
28.4	$C \leftarrow \text{linkageClusteribk} A$	476
28.5	$C \leftarrow \text{leaderCluster}_D^f(A)$	477
28.6	$C \leftarrow \text{leaderCluster}_D^a(A)$	477
29.1	$\text{distributedGraphColoring}$	494
29.2	$m_\varepsilon \leftarrow \text{errorBurst}(m, \varepsilon)$	500

List of Listings

4.1	Two examples for the PL dialect used by Cramer for Genetic Programming . .	141
4.2	An example for the JB Mapping	141
4.3	Another example for the JP Mapping	141
4.4	A trivial symbolic regression grammar.	148
4.5	A simple grammar for C functions that could be used in Gads.	150
4.6	A simple grammar for C functions that could be used by GE.	152
4.7	A Christiansen grammar for C functions that that use variables.	157
4.8	A complex conditional statement.	176
4.9	The RBGP version of Listing 4.8.	177
4.10	An equivalent alternative version of Listing 4.9.	177
4.11	A loop in Rule-based Genetic Programming.	177
4.12	The RBGP-version of Listing 4.11.	177
4.13	An equivalent alternative version of Listing 4.12.	177
4.14	A simple selection sort algorithm written in the eRBGP language.	178
4.15	The eRBGP version of Listing 4.8 and Listing 4.9.	179
4.16	The eRBGP version of Listing 4.11 and Listing 4.12.	179
4.17	A first, simple example for a Push program.	181
4.18	An example for the usage of the CODE stack.	181
4.19	Another example for the usage of the CODE stack.	181
4.20	An example for the creation of procedures.	182
4.21	An example for the creation of procedures similar to Listing 4.20.	182
20.1	An example Royal Road function.	286
20.2	Some training cases for the GCD problem.	307
20.3	The RBGP version of the Euclidean algorithm.	312
20.4	An overfitted RBGP solution to the GCP problem.	312
23.1	The optimal broadcasting scheme evolved for 2-dimensional, directed grids. . .	358
25.1	The enum EClasses with the possible DMC 2007 classifications.	397
25.2	The structure of our DMC 2007 classifier system.	399
25.3	The embryogeny component of our DMC 2007 contribution.	400
25.4	The simulation for testing the DMC 2007 classifier system.	401
25.5	The profit objective function $f_1(C) = -P(C)$ for the DMC 2007.	402
25.6	The size objective function $f_2(C) = C $ for the DMC 2007.	402
25.7	A main method that runs the evolution for the 2007 DMC.	404
27.1	Approximating D^2X of $r(y)$	465
29.1	A simple generative grammar.	502
29.2	An example context-free generative grammar G	503
29.3	An example expansion of G	504
29.4	Natural numbers – a small BNF example.	504
29.5	Integer numbers – a small EBNF example.	505

29.6 A simple context-free grammar.	506
29.7 A small example for attribute grammars.	506
29.8 Syntax of an Extended Attribute Grammar symbol.	507
29.9 The small example G_1 for Extended Attribute Grammars.	507
29.10A typical expansion of G_1	507
29.11An extended attribute grammar G_2 for binary numbers.....	508
29.12A Christiansen Grammar creating character strings.	509
29.13Christiansen Grammar for a simple programming language.....	509
29.14Another simple context-free grammar.	510
29.15A small Lisp-example: How to compute Fibonacci numbers.....	512