

Applied and Numerical Harmonic Analysis

Jonathan Cohen
Ahmed I. Zayed
Editors

Wavelets and Multiscale Analysis

Theory and Applications

 Birkhäuser

Applied and Numerical Harmonic Analysis

Series Editor

John J. Benedetto

University of Maryland

Editorial Advisory Board

Akram Aldroubi

Vanderbilt University

Jelena Kovačević

Carnegie Mellon University

Andrea Bertozzi

University of California at Los Angeles

Gitta Kutyniok

University of Osnabrück

Douglas Cochran

Arizona State University

Mauro Maggioni

Duke University

Hans G. Feichtinger

University of Vienna

Zuowei Shen

National University of Singapore

Christopher Heil

Georgia Institute of Technology

Thomas Strohmer

University of California at Davis

Stéphane Jaffard

University of Paris XII

Yang Wang

Michigan State University

Jonathan Cohen
Ahmed I. Zayed
Editors

Wavelets and Multiscale Analysis

Theory and Applications

Editors

Jonathan Cohen
Department of Mathematical Sciences
DePaul University
2320 North Kenmore Avenue
Chicago, IL 60614-3250, USA
jcohen@depaul.edu

Ahmed I. Zayed
Department of Mathematical Sciences
DePaul University
2320 North Kenmore Avenue
Chicago, IL 60614-3250, USA
azayed@depaul.edu

ISBN 978-0-8176-8094-7 e-ISBN 978-0-8176-8095-4

DOI 10.1007/978-0-8176-8095-4

Springer New York Dordrecht Heidelberg London

Library of Congress Control Number: 2011922755

Mathematics Subject Classification (2010): 42C15, 42C40, 41A58, 41A65, 43A15, 46E20, 47B37, 47B38, 52C22, 53C80, 62H20, 85A40

© Springer Science+Business Media, LLC 2011

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

www.birkhauser-science.com

ANHA Series Preface

The *Applied and Numerical Harmonic Analysis (ANHA)* book series aims to provide the engineering, mathematical, and scientific communities with significant developments in harmonic analysis, ranging from abstract harmonic analysis to basic applications. The title of the series reflects the importance of applications and numerical implementation, but richness and relevance of applications and implementation depend fundamentally on the structure and depth of theoretical underpinnings. Thus, from our point of view, the interleaving of theory and applications and their creative symbiotic evolution is axiomatic.

Harmonic analysis is a wellspring of ideas and applicability that has flourished, developed, and deepened over time within many disciplines and by means of creative cross-fertilization with diverse areas. The intricate and fundamental relationship between harmonic analysis and fields such as signal processing, partial differential equations (PDEs), and image processing is reflected in our state-of-the-art *ANHA* series.

Our vision of modern harmonic analysis includes mathematical areas such as wavelet theory, Banach algebras, classical Fourier analysis, time–frequency analysis, and fractal geometry, as well as the diverse topics that impinge on them.

For example, wavelet theory can be considered an appropriate tool to deal with some basic problems in digital signal processing, speech and image processing, geophysics, pattern recognition, biomedical engineering, and turbulence. These areas implement the latest technology from sampling methods on surfaces to fast algorithms and computer vision methods. The underlying mathematics of wavelet theory depends not only on classical Fourier analysis, but also on ideas from abstract harmonic analysis, including von Neumann algebras and the affine group. This leads to a study of the Heisenberg group and its relationship to Gabor systems, and of the metaplectic group for a meaningful interaction of signal decomposition methods. The unifying influence of wavelet theory in the aforementioned topics illustrates the justification for providing a means for centralizing and disseminating information from the broader, but still focused, area of harmonic analysis. This will be a key role of *ANHA*. We intend to publish the scope and interaction that such a host of issues demands.

Along with our commitment to publish mathematically significant works at the frontiers of harmonic analysis, we have a comparably strong commitment to publish

major advances in the following applicable topics in which harmonic analysis plays a substantial role:

<i>Antenna theory</i>	<i>Prediction theory</i>
<i>Biomedical signal processing</i>	<i>Radar applications</i>
<i>Digital signal processing</i>	<i>Sampling theory</i>
<i>Fast algorithms</i>	<i>Spectral estimation</i>
<i>Gabor theory and applications</i>	<i>Speech processing</i>
<i>Image processing</i>	<i>Time–frequency and</i>
<i>Numerical partial differential equations</i>	<i>time-scale analysis</i>
	<i>Wavelet theory</i>

The above point of view for the *ANHA* book series is inspired by the history of Fourier analysis itself, whose tentacles reach into so many fields.

In the last two centuries, Fourier analysis has had a major impact on the development of mathematics, on the understanding of many engineering and scientific phenomena, and on the solution of some of the most important problems in mathematics and the sciences. Historically, Fourier series were developed in the analysis of some of the classical PDEs of mathematical physics; these series were used to solve such equations. In order to understand Fourier series and the kinds of solutions they could represent, some of the most basic notions of analysis were defined, e.g., the concept of “function”. Since the coefficients of Fourier series are integrals, it is no surprise that Riemann integrals were conceived to deal with uniqueness properties of trigonometric series. Cantor’s set theory was also developed because of such uniqueness questions.

A basic problem in Fourier analysis is to show how complicated phenomena, such as sound waves, can be described in terms of elementary harmonics. There are two aspects of this problem: first, to find, or even define properly, the harmonics or spectrum of a given phenomenon, e.g., the spectroscopy problem in optics; second, to determine which phenomena can be constructed from given classes of harmonics, as done, e.g., by the mechanical synthesizers in tidal analysis.

Fourier analysis is also the natural setting for many other problems in engineering, mathematics, and the sciences. For example, Wiener’s Tauberian theorem in Fourier analysis not only characterizes the behavior of the prime numbers, but also provides the proper notion of spectrum for phenomena such as white light; this latter process leads to the Fourier analysis associated with correlation functions in filtering and prediction problems, and these problems, in turn, deal naturally with Hardy spaces in the theory of complex variables.

Nowadays, some of the theory of PDEs has given way to the study of Fourier integral operators. Problems in antenna theory are studied in terms of unimodular trigonometric polynomials. Applications of Fourier analysis abound in signal processing, whether with the fast Fourier transform (FFT), or filter design, or the adaptive modeling inherent in time–frequency-scale methods such as wavelet theory. The coherent states of mathematical physics are translated and modulated

Fourier transforms, and these are used, in conjunction with the uncertainty principle, for dealing with signal reconstruction in communications theory. We are back to the *raison d'être* of the *ANHA* series!

University of Maryland
College Park

John J. Benedetto
Series Editor

Contents

1	An Introduction to Wavelets and Multiscale Analysis: Theory and Applications.....	1
	Ahmed I. Zayed	

Part I The Mathematical Theory of Wavelets

2	The Construction of Wavelet Sets.....	17
	John J. Benedetto and Robert L. Benedetto	
3	The Measure of the Closure of a Wavelet Set May Be $> 2\pi$	57
	Zhihua Zhang	
4	Quincunx Wavelets on \mathbb{T}^2	65
	Kenneth R. Hoover and Brody Dylan Johnson	
5	Crystallographic Haar-Type Composite Dilation Wavelets	83
	Jeffrey D. Blanchard and Kyle R. Steffen	
6	From Full Rank Subdivision Schemes to Multichannel Wavelets: A Constructive Approach.....	109
	Costanza Conti and Mariantonia Cotronei	
7	Unitary Systems and Bessel Generator Multipliers.....	131
	Deguang Han and David R. Larson	
8	The Zak Transform(s).....	151
	Eugenio Hernández, Hrvoje Šikić, Guido L. Weiss, and Edward N. Wilson	

Part II The Geometry of Large Data Sets

9	Harmonic Analysis of Digital Data Bases	161
	Ronald R. Coifman and Matan Gavish	

10	Some Recent Advances in Multiscale Geometric Analysis of Point Clouds	199
	Guangliang Chen, Anna V. Little, Mauro Maggioni, and Lorenzo Rosasco	
11	Signal Ensemble Classification Using Low-Dimensional Embeddings and Earth Mover's Distance	227
	Linh Lieu and Naoki Saito	
Part III Applications of Wavelets		
12	Wavelets on Manifolds and Statistical Applications to Cosmology	259
	Daryl Geller and Azita Mayeli	
13	Wavelets, a Numerical Tool for Atmospheric Data Analysis	279
	Parick Fischer and Ka-Kit Tung	
14	Denoising Speech Signals for Digital Hearing Aids: A Wavelet Based Approach	299
	Nathaniel Whitmal, Janet Rutledge, and Jonathan Cohen	
	Index	333

Preface

This book is a collection of papers written or co-authored by participants in the “Twenty Years of Wavelets” conference held at DePaul in May, 2009. The conference attracted almost a hundred participants from five different countries over three days. There were 13 plenary lectures and 16 contributed talks. The conference was envisioned to celebrate the twentieth anniversary of a one-day conference on applied and computational harmonic analysis held at DePaul in May 1989 and was organized by one of the editors, Jonathan Cohen. The 1989 DePaul conference was scheduled to supplement a two-day special session of a regional AMS meeting on computational harmonic analysis and approximation theory. Combined together, the three days of talks may have been the first conference in the United States which featured the subject of wavelets. Although the focus of that conference was computational harmonic analysis, wavelet theory, which was in its infancy at the time, played a central role in the three days of talks.

After two decades of extensive research activities, it was appropriate to pause and have a look back at what had been accomplished and ponder what lay ahead. This was exactly the aim of the 2009 conference. The conference had two sub-themes, past and future. Some of the plenary speakers, including I. Daubechies and J. Kovačević, gave expository and survey talks covering the history and major accomplishments in the field and some speakers focused on new directions for wavelets, especially in the area of geometric harmonic analysis.

All conference speakers were invited to submit papers related to the themes of the conference. This was interpreted broadly to include articles in applied and computational harmonic analysis. Though many of the articles are based on conference presentations, this book was not envisioned as a proceedings and some of the articles represent material not presented at the conference. All the papers in this book were anonymously refereed.

The book is divided into three parts. The first is devoted to the mathematical theory of wavelets and features several papers on the geometry of sets and the development of wavelet bases. The second part deals with the underlying geometry of large data sets and how tools of harmonic analysis prove useful in extracting information from them. The third part is devoted to exploring some ways that harmonic analysis, and wavelet theory in particular, have been applied to study real-world problems.

The articles in this book are mostly written by mathematicians and are intended for mathematicians and engineers with some background in Fourier analysis and the theory of wavelets. The book should be accessible to workers in the field and to graduate students with an interest in working in related areas.

We gratefully acknowledge the National Science Foundation, NSF Grant DMS-0852170, and DePaul University, both of whom provided generous financial support for the conference. We would also like to express our appreciation to the authors for submitting their work and meeting the deadlines, to the referees for their help and cooperation, and to Thomas Grasso, the science editor for Birkhäuser, for his support throughout this project.

Finally, we note with sadness that one of the authors in this volume, Daryl Geller, passed away in late January. He was a very fine mathematician who will be missed by his colleagues and friends.

DePaul University, Chicago, Illinois
February, 2011

Jonathan Cohen
Ahmed I. Zayed

Contributors

John J. Benedetto Department of Mathematics, University of Maryland, College Park, MD 20742, USA, jjb@math.umd.edu

Robert L. Benedetto Department of Mathematics, Amherst College, Amherst, MA 01002, USA, rlb@math.amherst.edu

Jeffrey D. Blanchard Department of Mathematics and Statistics, Grinnell College, Grinnell, Iowa, USA, jeff@math.grinnell.edu

Guangliang Chen Duke University, Durham, NC, USA, glchen@math.duke.edu

Jonathan Cohen Department of Mathematics, DePaul University, Chicago, Illinois, USA, jcohen@math.depaul.edu

Ronald R. Coifman Program in Applied Mathematics, Yale University, 51 Prospect Street, New Haven, CT 06511, USA, coifman-ronald@yale.edu

Costanza Conti Dipartimento di Energetica “Sergio Stecco”, Università di Firenze, Via Lombroso 6/17, I-50134 Firenze, Italy, costanza.conti@unifi.it

Mariantonio Cotronei DIMET, Università “Mediterranea” di Reggio Calabria, Via Graziella, I-89122 Reggio Calabria, Italy, mariantonia.cotronei@unirc.it

Patrick Fischer Institut de Mathématiques de Bordeaux, Université Bordeaux 1, 33405 Talence Cedex, France, patrick.fischer@math.u-bordeaux1.fr

Matan Gavish Department of Statistics, Stanford University, Sequoia Hall, Stanford, CA 94305, USA, gavish@stanford.edu
and
Program in Applied Mathematics, Yale University, New Haven, CT 06511, USA

Daryl Geller Mathematics Department, Stony Brook University, Stony Brook, NY 11794, USA

Deguang Han Department of Mathematics, University of Central Florida, Orlando, FL 32816, USA, ghan@pegasus.cc.ucf.edu

Eugenio Hernández Departamento de Matemáticas, Universidad Autónoma de Madrid, 28049 Madrid, Spain, eugenio.hernandez@uam.es

Kenneth R. Hoover California State University Stanislaus, Turlock, CA, USA,
khoover@csustan.edu

Brody Dylan Johnson Saint Louis University, St. Louis, MO, USA,
brody@slu.edu

David R. Larson Department of Mathematics, Texas A&M University, College Station, TX, USA, larson@math.tamu.edu

Linh Lieu Department of Mathematics, University of California, One Shields Avenue, Davis, CA 95616, USA, llieu@math.ucdavis.edu

Anna V. Little Duke University, Durham, NC, USA, avl@math.duke.edu

Mauro Maggioni Duke University, Durham, NC, USA, mauro@math.duke.edu

Azita Mayeli Mathematics Department, City College of Technology, CUNY, 300 Jay Street, Brooklyn, NY 11201, USA, amayeli@citytech.cuny.edu

Lorenzo Rosasco Massachusetts Institute of Technology, USA, lrosasco@mit.edu
and
Istituto Italiano di Tecnologia, Italy, CBCL

Janet Rutledge Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County, USA, jrutledge@umbc.edu

Naoki Saito Department of Mathematics, University of California, One Shields Avenue, Davis, CA 95616, USA, saito@math.ucdavis.edu

Hrvoje Šikić Department of Mathematics, University of Zagreb, Bijenička 30, HR-10000, Zagreb, Croatia, hsikic@math.hr

Kyle R. Steffen Department of Mathematics, University of Utah, Salt Lake City, Utah, USA, kyle.steffen@gmail.com

Ka-Kit Tung Applied Mathematics Department, University of Washington, Seattle, WA 98195-2420, USA, tung@math.washington.edu

Guido L. Weiss Department of Mathematics, Washington University, Campus Box 1146, St. Louis, MO 63130, USA, guido@math.wustl.edu

Nathaniel Whitmal Department of Communication Disorders, University of Massachusetts, Amherst, Amherst, Massachusetts, USA,
nwhitmal@comdis.umass.edu

Edward N. Wilson Department of Mathematics, Washington University, Campus Box 1146, St. Louis, MO 63130, USA, enwilson@math.wustl.edu

Ahmed I. Zayed Department of Mathematical Sciences, DePaul University, Chicago, IL 60614, USA, azayed@condor.depaul.edu

Zhihua Zhang College of Global Change and Earth System Science, Beijing Normal University, Beijing 100875, China
and

Department of Mathematics, University of California, Davis, CA 95616, USA,
zhangzh@math.ucdavis.edu

Chapter 1

An Introduction to Wavelets and Multiscale Analysis: Theory and Applications

Ahmed I. Zayed

Abstract The purpose of this introductory chapter is to give the reader an overview of the contents of the monograph and show how the chapters are tied together. We give a brief description of each chapter but with emphasis on how the chapters fit in the monograph and the general subject area. The descriptions are not meant to replace, but to supplement, the chapters' abstracts, which summarize the chapters' main results.

1.1 Introduction

This monograph is broadly based on talks given at an international conference on wavelets that was held at DePaul University, May 15–17, 2009, and was partially supported by grants from the National Science Foundation and DePaul University Research Council. The title of the conference was “Twenty Years of Wavelets” to commemorate the twentieth anniversary of another conference on wavelets that was held at the same university in 1989.

Since the introduction of wavelets in the early 1980s, the subject has undergone tremendous developments both on the theoretical and applied fronts. Myriads of research and survey papers and monographs have been published on the subject covering different areas of applications, such as signal and image processing, denoising, and data compression. This monograph not only contributes to this burgeoning subject, but also sheds light on new directions for wavelets, especially in the area of *geometric harmonic analysis* which aims at developing harmonic analysis techniques to deal with large data sets in high dimensions. This approach, which was pioneered by R. Coifman and his team at Yale University, has shown very promising results.

A.I. Zayed (✉)

Department of Mathematical Sciences, DePaul University, Chicago, IL 60614, USA

e-mail: azayed@condor.depaul.edu

R. Coifman and some of his collaborators and former students have contributed chapters to this book. These chapters may provide researchers and graduate students an opportunity to learn about recent developments in the area of multiscale harmonic analysis.

The purpose of this chapter is to give the reader an overview of the contents of the monograph and show how the chapters are connected together. We assume that the reader is familiar with the rudiments of wavelets and multiresolution analyses. The book is divided into three parts: the first part is the Mathematical Theory of Wavelets, the second is the Multiscale Analysis of Large Data Sets, and the third is Applications of Wavelets. The chapters in the first part are grouped together by common themes: wavelet sets and wavelet construction. Chapters 2 and 3 deal with wavelet sets, while Chaps. 4–6 discuss the construction of wavelets in different settings, such as wavelets on a torus, crystallographic composite dilation wavelets, and vector-valued (multichannel) wavelets. The second part comprises chapters on multiscale analysis of large data sets. The chapters in part three discuss applications of wavelets in three different fields: cosmology, atmospheric data analysis, and denoising speech signals for digital hearing aids.

Admittedly, the boundaries between the three parts are rather subjective. Chapter 12 would have also fit nicely in the first part of the book because it introduces a wavelet construction on compact Riemannian manifolds; however, it is placed in the third part because of the authors' emphasis on the applications of their work to cosmology.

The notation we use in this chapter is standard. We denote the sets of real numbers by \mathbb{R} , the integers by \mathbb{Z} , and the natural numbers by \mathbb{N} . Recall that the standard dyadic wavelets are functions of the form

$$\psi_{m,n}(x) = 2^{n/2} \psi(2^n x - m), m, n \in \mathbb{Z}, x \in \mathbb{R}, \quad (1.1)$$

that are generated from one single function ψ , called the mother wavelet, by translation and dilation. More generally, if we denote the translation, dilation, and modulation operators, respectively, by T_m, D^n, E_k where

$$T_m f(x) = f(x - m), \quad D^n f(x) = 2^{n/2} f(2^n x), \quad E_k f(x) = e^{i\langle k, x \rangle} f(x),$$

where $x \in \mathbb{R}^d, m, k \in \mathbb{Z}^d, n \in \mathbb{Z}$, then the wavelet and Gabor systems on \mathbb{R}^d can be written, respectively, as

$$\psi_{m,n}(x) = D^n T_m \psi(x), \quad g_{m,n}(x) = E_k T_m g(x).$$

A more general form of wavelets in \mathbb{R}^d is given by

$$(\psi_j^A)_{m,n}(x) = |\det A|^{n/2} \psi_j(A^n x - m), \quad (1.2)$$

where A is a real expansive $d \times d$ invertible matrix and $j = 1, \dots, J, n \in \mathbb{Z}, m \in \mathbb{Z}^d$.

A multiresolution analysis of $L^2(\mathbb{R})$ consists of a nested sequence of closed subspaces $\{V_n\}_{n=-\infty}^{\infty}$ of $L^2(\mathbb{R})$ and a function $\phi \in V_0$, called the scaling function, such that

- (1) $\dots \subset V_{-1} \subset V_0 \subset V_1 \subset V_2 \subset \dots$.
- (2) $\bigcup_{i=-\infty}^{\infty} V_i$ is dense in $L^2(\mathbb{R})$.
- (3) $\bigcap_{i=-\infty}^{\infty} V_i = \{0\}$.
- (4) $f(x) \in V_0 \Leftrightarrow f(2^j x) \in V_j$.
- (5) $\{T_n \phi(x) = \phi(x - n)\}_{n=-\infty}^{\infty}$ is an orthonormal (Riesz) basis for V_0 .

It is well known that given a multiresolution analysis, one can construct an orthonormal wavelet basis of $L^2(\mathbb{R})$.

1.1.1 Mathematical Theory of Wavelets

The construction of dyadic orthonormal wavelet bases in $L^2(\mathbb{R})$ hinges on the construction of the mother wavelet ψ . The pioneering work of Y. Meyer and S. Mallat [5, 6] gave an algorithm for constructing the mother wavelet in the setting of multiresolution analysis (MRA). However, not every wavelet is generated from a multiresolution analysis as J. Journé in 1992 demonstrated by his celebrated example of a non-MRA wavelet basis for $L^2(\mathbb{R})$; see [3, p. 136]. In higher dimensions, the construction of orthonormal wavelet bases was more elusive. The most common construction of an orthonormal wavelet basis came from the theory of multiresolution analysis (MRA) which requires $2^d - 1$ functions $\psi_j, j = 1, \dots, 2^d - 1$, to generate the resulting orthonormal basis (ONB), $(\psi_j)_{m,n}$, of $L^2(\mathbb{R}^d)$; see [6, p. 90].

For some time there was doubt about the existence of a single dyadic orthonormal wavelet basis for $\mathbb{R}^d, d > 1$. The work of Dai, Larson, and Speegle [1, 2] on operator theory proved the existence of such wavelets in $L^2(\mathbb{R}^d), d > 1$. Their proof used the notion of wavelet sets and operator algebra methods. It turned out that the Fourier transform of such a mother wavelet ψ is the characteristic function χ_{Ω} of a measurable set Ω , which is called a wavelet set.

The construction of such a wavelet set is not obvious, but it is necessary that its translates by \mathbb{Z}^d provide a tiling of \mathbb{R}^d . In fact, it is known that if $E \subset \mathbb{R}$ is a measurable set, then E is a wavelet set if and only if E is both a 2-dilation generator (modulo null sets) of a partition of \mathbb{R} and a 2π -translation generator (modulo null sets) of a partition of \mathbb{R} . The simplest example of a wavelet set is given by the Shannon's wavelet ψ whose Fourier transform $\hat{\psi} = \chi_E$, where $E = [-2\pi, -\pi) \cup [\pi, 2\pi)$. The structure of wavelet sets in higher dimensions is much more intricate.

In Chap. 2, John Benedetto and his son Robert give a general method for constructing single dyadic wavelets, which generate wavelet orthonormal bases (ONBs) for the space of square-integrable functions in two important antipodal cases. These cases are $L^2(\mathbb{R}^d)$, where \mathbb{R}^d is the d -dimensional Euclidean space, and $L^2(G)$, where G belongs to the class of locally compact Abelian groups (LCAGs) which contain a compact open subgroup. The wavelets they construct are not derived from any MRA.

In the first five sections, the authors discuss the geometry and construction of Euclidean wavelet sets and then proceed to extend their results to the non-Euclidean cases, such as locally compact Abelian groups and p -adic fields, \mathbb{Q}_p , of p -adic rationals. One of the salient features of this well written chapter is that it is almost self contained. The authors give a historical account of the subject, as well as the needed definitions and mathematical tools.

It is a well known fact that if ψ is an orthonormal wavelet, i.e., it generates an orthonormal wavelet basis in $L^2(\mathbb{R})$, then $\sum_k |\hat{\psi}(\omega + 2\pi k)|^2 = 1$, and hence $|\hat{\psi}(\omega)| \leq 1$. Furthermore,

$$2\pi = \int_{\mathbb{R}} |\hat{\psi}(\omega)|^2 d\omega = \int_E |\hat{\psi}(\omega)|^2 d\omega \leq |E| \leq |\bar{E}| = |\text{supp } \hat{\psi}|,$$

where $E = \{\omega : \hat{\psi}(\omega) \neq 0\}$ and \bar{E} is the closure of E . Thus, the minimal measure of the support of the Fourier transform of a wavelet ψ is 2π and clearly the support of the Fourier transform of ψ possesses the minimal measure if and only if

$$|\hat{\psi}| = \chi_E \text{ and } |E| = |\bar{E}| = 2\pi,$$

where χ_E is the characteristic function of E . Recall that in the latter case, where the Fourier transform of ψ is the a characteristic function of a set E , the set E is called a wavelet set. Many wavelet sets have been constructed where E is a finite union of intervals. Of course, in these cases $|E| = |\bar{E}| = 2\pi$ and even in the known cases where E is a union of infinitely many intervals, still $|E| = |\bar{E}| = 2\pi$. Hence, this raises the question of whether there exists a wavelet set E^* such that $|\bar{E}^*| > |E^*| = 2\pi$. In Chap. 3, Z. Zhang answers this question in the affirmative. He also answers the same question for the scaling function of the associated multiresolution analysis.

One of the key elements in the construction of a multiresolution analysis is the notion of shift-invariant spaces and the nested sequence of subspaces generated from them by dilation. These ideas are extended by K. Hoover and B. Johnson in Chap. 4 to construct a multiresolution analysis and hence a finite-dimensional system of orthonormal wavelets in $L^2(\mathbb{T})$, where \mathbb{T} is the torus. In this setting, dilation and translation have different meanings. The dilation operation is achieved through a matrix A , called the Quincunx dilation matrix

$$A = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix};$$

see (1.2), while translation is considered over a discrete subgroup of \mathbb{T}^2 . More precisely, for a fixed integer $j > 0$, a lattice, Γ_j of order 2^j generated by A is a collection of 2^j distinct coset representatives of $A^{-j}\mathbb{Z}^2/\mathbb{Z}^2$. The translation operators are generated by elements of Γ_j . A shift-invariant space is defined as a space that consists of functions in $L^2(\mathbb{T}^2)$ that are invariant under translation by elements of Γ_j . Having introduced the basic tools for shift-invariant spaces, the authors go on to construct a MRA of order 2^j consisting of closed subspaces $\{V_k\}_{k=0}^j$ of $L^2(\mathbb{T}^2)$ satisfying

similar properties to those of the standard MRA, such as the nested sequence property and the existence of a scaling function. They conclude their work by giving examples of wavelet systems on the torus that are analogs of the Shannon and Haar wavelets.

Another type of wavelets, called crystallographic Haar-type composite dilation wavelets, is discussed in Chap. 5 by J. Blanchard and K. Steffen. These wavelets are composite dilation wavelets which arise from crystallographic groups and are linear combinations of characteristic functions. To briefly explain some of the terminology, let $GL_d(\mathbb{R})$ be the group of $d \times d$ invertible matrices. For any $A \in GL_d(\mathbb{R})$, define the dilation by (1.2). A full rank lattice Γ is a subset of \mathbb{R}^d with the property that there exists $A \in GL_d(\mathbb{R})$, such that $\Gamma = A\mathbb{Z}^d$. A group of invertible matrices G and a full rank lattice Γ are said to satisfy the crystallographic condition if Γ is invariant under the action of G , i.e. $G(\Gamma) = \Gamma$. The translation of f by $k \in \Gamma$ is defined as usual by $T_k f(x) = f(x - k)$. With these two unitary operators, an affine system,

$$\mathcal{U}_{C,\Gamma} = \left\{ D_c T_k \psi^\ell(x) : c \in C, k \in \Gamma, \ell = 1, \dots, L \right\},$$

is constructed from a countable set of invertible matrices, $C \subset GL_d(\mathbb{R})$, a full rank lattice Γ , and a set of generating functions, $\Psi(x) = (\psi^1, \dots, \psi^L) \subset L^2(\mathbb{R}^d)$. An affine system with composite dilations are obtained when $C = AB$ is the product of two subsets of invertible matrices A and B .

In this chapter, it is assumed that $A = \{a^j : j \in \mathbb{Z}\}$ is a group generated by integer powers of an expanding matrix, a , and B is a subgroup of invertible matrices so that the affine system is in the form

$$\mathcal{U}_{a,B,\Gamma}(\Psi) = \left\{ D_a^j D_b T_k \psi^\ell : j \in \mathbb{Z}, b \in B, k \in \Gamma, 1 \leq \ell \leq L \right\}.$$

The system of functions $\Psi = (\psi^1, \dots, \psi^L) \subset L^2(\mathbb{R}^d)$ is called a composite dilation wavelet if $\mathcal{U}_{a,B,\Gamma}(\Psi)$ is an orthonormal basis of $L^2(\mathbb{R}^d)$.

When $B = I_d$, where I_d is the identity matrix, we obtain the standard multiwavelet definition, and when $d = 1, \Gamma = \mathbb{Z}$, and $a = 2$, we obtain the standard dyadic wavelets. The authors discuss examples of Haar-type composite dilation wavelets that were introduced in [4] by Krishtal, Robinson, Weiss, and Wilson under the assumptions: (a) B is a finite group and (b) the lattice Γ is invariant under the action of B , i.e., $B(\Gamma) = \Gamma$. It is shown that assumption (b) implies assumption (a), i.e., B must be a finite group in a Haar-type composite dilation wavelet system. Assumption (b) is the crystallographic condition. The authors develop a systematic way to construct crystallographic Haar-type composite dilation wavelets for $L^2(\mathbb{R}^d)$ and then discuss in more details, with examples, the construction for $L^2(\mathbb{R}^2)$.

Chapter 6 by C. Conti and M. Cotronei deals with the construction of wavelets for the analysis of vector-valued functions. Such functions arise naturally in many applications where the data to be processed are samples of vector-valued functions. Multichannel signals, or vector-valued signals whose components come from different sources with possible intrinsic correlations, such as brain activity (EEG/MEG) data or colored images, may exhibit a high correlation which can be revealed and exploited by what is called multichannel wavelet analysis.

The most common signal processing technique to handle vector-valued data is to deal with the signal's components one by one and thus ignoring possible relationships between some components. The standard scalar wavelet analysis, including multiwavelets, does not take into account the correlations among components. Recently, matrix wavelets, multichannel wavelets (MCW), and multichannel multiresolution analyses (MCMRA) have been proposed for the analysis of matrix-valued and vector-valued signals. An underlying concept in these different schemes is the existence of a matrix refinable function which satisfies the so-called full rank condition.

To explain some of the terminology, let $r, s \in \mathbb{N}$, and $A \in \mathbb{R}^{r \times s}$ be an $r \times s$ matrix. Let $\mathcal{A} = (A_j \in \mathbb{R}^{r \times s}, j \in \mathbb{Z})$ be a bi-infinite sequence of such matrices such that

$$\|\mathcal{A}\|_2 = \left(\sum_{j \in \mathbb{Z}} |A_j|_2^2 \right)^{1/2} < \infty,$$

where $|\cdot|_2$ denotes the standard ℓ^2 norm of $r \times s$ matrices. Let $L_2^{r \times s}(\mathbb{R})$ denote the Banach space of $r \times s$ matrix-valued functions on \mathbb{R} with components in $L^2(\mathbb{R})$ and norm

$$\|F\|_2 = \left(\sum_{k=1}^s \sum_{j=1}^r \int_{\mathbb{R}} |F_{j,k}(x)|^2 dx \right)^{1/2}.$$

For a matrix function F and a matrix sequence \mathcal{A} , the convolution $*$ is defined as

$$F * \mathcal{A} = \sum_{k \in \mathbb{Z}} F(\cdot - k) A_k.$$

Fix a matrix sequence $\mathcal{A} = (A_j \in \mathbb{R}^{r \times r}, j \in \mathbb{Z})$ and for any bi-infinite vector sequence $c = (c_j \in \mathbb{R}^r, j \in \mathbb{Z})$ define the vector subdivision operator $S_{\mathcal{A}}$ based on the matrix mask \mathcal{A} , as

$$S_{\mathcal{A}}(c) = \left(\sum_{k \in \mathbb{Z}} A_{j-2k} c_k, j \in \mathbb{Z} \right).$$

A vector subdivision scheme is defined inductively as

$$c^0 = c, \quad c^n = S_{\mathcal{A}}(c^{n-1}).$$

A multichannel multiresolution analysis (MCMRA) in the space $L_2^r(\mathbb{R})$ of square integrable vector-valued functions can be defined as a nested sequence $\cdots V_{-1} \subset V_0 \subset V_1 \subset \cdots$ of closed subspaces of $L_2^r(\mathbb{R})$ with similar properties to those in the scalar case, but in which the role of the scaling function ϕ is now replaced by a vector function so that the space V_0 is generated by the integer translates of r function vectors, f^i . That is, there exist $f^i = (f_1^i, \dots, f_r^i) \in L_2^r(\mathbb{R}), i = 1, \dots, r$, so that any $h \in V_0$ can be written as $h = F * c$, where $F = (f^1, \dots, f^r) \in L_2^{r \times r}(\mathbb{R})$ and $c = (c_k)$. The matrix F is called the matrix scaling function and it is required that F be stable, i.e., there exist K_1, K_2 such that

$$K_1 \|c\|_2 \leq \|F * c\|_2 \leq K_2 \|c\|_2.$$

It is shown that the subspaces V_j form a MCMRA if there exists a full rank stable matrix refinable function $F \in L_2^{r \times r}(\mathbb{R})$ such that

$$V_j = \{F * c(2^j \cdot)\}, j \in \mathbb{Z}.$$

The reader who is familiar with MRA can now see the analogy with the scalar case and how one can associate a matrix wavelet to any orthogonal MCMRA.

In addition, the authors investigate full rank interpolatory schemes and show their connection to matrix refinable functions and multichannel wavelets. They show how to solve matrix quadrature mirror filter equations and give a constructive scheme that uses spectral factorization techniques and a matrix completion algorithm based on the solution of generalized Bezout identities.

In Chap. 7, D. Han and D. Larson use operator algebra techniques to study, in a unified way, wavelet and Gabor frames, or more generally, Bessel wavelets and Gabor–Bessel sets.

Recall that a sequence of vectors $\{x_n\}$ in a separable infinite dimensional Hilbert space \mathcal{H} is said to be a frame if there exist $A, B > 0$ such that for all $x \in \mathcal{H}$, we have

$$A \|x\|^2 \leq \sum_n |\langle x, x_n \rangle|^2 \leq B \|x\|^2.$$

If only the right-hand inequality holds, the sequence is called a Bessel Sequence. A unitary system \mathcal{U} is a set of unitary operators containing the identity operator I and acting on the Hilbert space \mathcal{H} . A vector $x \in \mathcal{H}$ is called complete wandering vector (respectively frame generator vector, or Bessel generator vector) for \mathcal{U} if

$$\mathcal{U}x = \{Ux : U \in \mathcal{U}\}$$

is an orthonormal basis (respectively a frame, or a Bessel sequence). A bounded linear operator A on \mathcal{H} that maps every frame (Bessel) generator for \mathcal{U} to a frame (Bessel) generator for \mathcal{U} is called a frame (Bessel) generator multiplier.

In this chapter, the authors characterize some special Bessel generator multipliers for unitary systems that are ordered products of two unitary groups. This includes the wavelet and the Gabor unitary systems since wavelets can be viewed as being generated by an ordered product $\{D^n T_m\}$ of two unitary groups, the dilation group $\{D^n\}$ and the translation group $\{T_k\}$, and acting on a separable infinite dimensional Hilbert space $L^2(\mathbb{R}^d)$. In the Gabor theory, the groups are the modulation group $\{E_k\}$ and the translation group $\{T_k\}$.

The authors point out that one can gain much additional perspective on wavelets and frame wavelets if one views them as special cases of Bessel wavelets. One reason for this is that, unlike the sets of wavelets and frame wavelets, the set of Bessel wavelets is a linear space. Some new results for wavelets and frame wavelets, as well as, some new proofs of previously known results, are obtained as special cases of the Bessel wavelet case. The same can be said about Gabor–Bessel generators. The chapter provides a detailed exposition of part of the history leading up to this work.

With the advent of wavelets and multiresolution analyses, shift-invariant spaces have become the focus of many research papers. Sampling spaces are special cases of shift-invariant spaces and they may be defined as

$$V = \left\{ f : f \in L^2(\mathbb{R}), \quad f(x) = \sum_k f(t_k) \phi(x - t_k), \quad \phi \in L^2(\mathbb{R}) \right\},$$

where $\{t_k\}$ is an increasing sequence of real numbers such that $\lim_{k \rightarrow -\infty} t_k = -\infty$, $\lim_{k \rightarrow \infty} t_k = \infty$, and $0 < \delta \leq |t_{k+1} - t_k|$, and $\{\phi(x - t_k)\}$ is a frame, or Riesz basis, or an orthonormal basis in $L^2(\mathbb{R})$. In most cases of interest, $t_k = k \in \mathbb{Z}$, and in this case the prototype is the Paley–Wiener space of functions f bandlimited to $[-\pi, \pi]$ and which have the representation

$$f(x) = \sum_k f(k) \frac{\sin \pi(x - k)}{\pi(x - k)} = \sum_k f(k) \phi(x - k),$$

where $\phi(x) = \text{sinc } x = \sin \pi x / \pi x$.

An important tool in the study of sampling spaces is the Zak transform, which is defined for $f \in L^2(\mathbb{R})$ by

$$(Zf)(x, w) = \sum_{k \in \mathbb{Z}} f(x + k) e^{-2\pi i k w}, \quad f \in L^2(\mathbb{R}). \quad (1.3)$$

In the last chapter of Part I, E. Hernández, H. Šikić, G. Weiss, and E. Wilson give a new insight into the Zak transform and its properties. Let $\phi(x, w)$ denote the function given by (1.3) and M be the space of all such functions. It is easy to see that ϕ is periodic in w with period one and that

$$\phi(x + \ell, w) = e^{2\pi i \ell w} \phi(x, w), \quad \ell \in \mathbb{Z}.$$

Moreover, it is known that $\phi \in L^2(T^2)$, where $T = [-1/2, 1/2)$ and that the linear operator Z maps $L^2(\mathbb{R})$ isometrically onto M .

Let us define $\tilde{\phi} \in L^2(T^2)$ by

$$\tilde{\phi}(x, w) = \sum_{k \in \mathbb{Z}} f(w + k) e^{2\pi i k x},$$

and denote the space of all such functions by \tilde{M} . It is easy to see that $\tilde{\phi}$ is periodic in x with period one and $\tilde{\phi}(x, w + \ell) = e^{-2\pi i \ell x} \tilde{\phi}(x, w)$. The authors introduce the unitary map

$$(U\phi)(x, w) = e^{-2\pi i x w} \phi(x, w)$$

on M , and also the Zak-like transform \tilde{Z} by

$$(\tilde{Z}g)(x, w) = \sum_{\ell \in \mathbb{Z}} g(w + \ell) e^{2\pi i \ell x} = \tilde{\phi}(x, w), \quad g \in L^2(\mathbb{R}).$$

It is shown that

$$\tilde{Z}^{-1}UZf = \hat{f}, \quad Z^{-1}U^*\tilde{Z}g = \check{g},$$

where \hat{f} denotes the Fourier transform of f and \check{g} denotes the inverse Fourier transform of g . The authors also show how to use the Zak transform to obtain elementary proofs of some results in harmonic analysis, including the Plancherel theorem and the Shannon sampling theorem.

1.1.2 Multiscale Analysis of Large Data Sets

R. Coifman and M. Gavish in Chap. 9 introduce digital data bases represented by data matrices M using harmonic analysis techniques. The prototype matrix M , is a data matrix whose set of columns X may be interpreted as observations or data points and set of rows Y may be interpreted as variables measured on the data points. The classical tools of multivariate statistics do not apply well in this case because one of the basic assumptions in multivariate statistics, namely, that the observations are independent and identically distributed, is no longer valid because, in general, correlations exist among the rows and columns. In this work, the authors propose another approach using harmonic analysis technique.

The authors first introduce the notion of a geometry and a Haar-like orthonormal basis on an abstract set, such as a set of observations. They then proceed to construct two coupled geometries on the rows and columns of the data matrices, which is called the coupled geometry of the matrix M . The coupled geometry is constructed using what the authors call a partition tree.

Guided by a classical result of J. O. Strömberg [7] that the tensor product of Haar bases is very efficient in representing functions on product spaces, the authors construct a Haar-like basis on X and Y , denoted by ψ_i and ϕ_j , and then show that the tensor basis $\{\psi_i \times \phi_j\}$ is an efficient basis of M . The Haar-like basis is a multiscale, localized orthonormal basis induced by a partition tree. Once a coupled geometry that is compatible with the matrix is obtained, the data matrix is expanded in the tensor-Haar basis and the data set is processed in the coefficient domain. It is shown that ℓ_p entropy conditions on the expansion coefficients of the database, viewed as a function on the product of the geometries, imply both smoothness and efficient reconstruction.

The authors describe how a tensor-Haar basis, induced by a coupled geometry that is compatible with the data matrix, can be used to perform compression, as well as, statistical tasks such as denoising and learning on the data matrix. They illustrate their technique for finding the coupled geometry of matrix rows and columns by giving two examples; the first example involves potential interaction between two point clouds in three dimensions and the second involves organizing a matrix whose rows are samples of some manifold Laplacian eigenfunctions.

As an interesting concrete example of data analysis using the general procedure described above, the authors consider a term-document data matrix. The documents

are abstracts of 1,047 articles obtained from the Science News journal website of different fields. One thousand words with the highest correlation to article subject classification are selected so that the $M_{i,j}$ entry of the data matrix M is the relative frequency of the i th word in the j th document, $i = 1, \dots, 1,000$ and $j = 1, \dots, 1,024$.

Chapter 10 by G. Chen, A. Little, M. Maggioni, and L. Rosasco, reports on some advances in multiscale geometric analysis for the study of large data sets that lie in high dimensional spaces but are confined to low-dimensional nonlinear manifolds. Symbolically, the data sets lie on a d -dimensional nonlinear manifold M embedded in a high dimensional space \mathbb{R}^D and are corrupted by high-dimensional noise.

The data sets are represented as discrete sets in \mathbb{R}^D , where D can be as large as 10^6 . Unlike classical statistics, which deals with large data sets of size n in a space of fixed dimension D , with at least $n \gg 2^D$, in the case under consideration here we have n of the same order as D , and oftentimes $n < D$.

One of the statistical tools used to analyze data sets is the Principle Component Analysis (PCA) which is based on the Singular Value Decomposition (SVD). Recall that if X is an $n \times D$ matrix, it can be decomposed as $X = UYV^t$, where $U \in \mathbb{R}^{n \times n}$ and $V \in \mathbb{R}^{D \times D}$ are orthogonal and $Y \in \mathbb{R}^{n \times D}$, is diagonal and semi positivedefinite. The diagonal elements $\{\lambda_i\}$ of Y are called singular values and are ordered in decreasing order of magnitude and called the SVD of X . The first d columns of V provide the d -dimensional least square fit to X . If the rows $\{x_i\}$ of X represent n data points in \mathbb{R}^D , and lie on a bounded domain in a d -dimensional linear subspace of \mathbb{R}^D , then the SVD method may be applied.

However, in this chapter, the authors focus on the case where the data lie on a *non-linear* d -dimensional manifold \mathcal{M} in \mathbb{R}^D with $d \ll D$, and where the SVD method does not work. They propose a new technique based on what they call *geometric wavelets* which aims at efficiently representing the data. Essentially, the authors construct a data-dependent dictionary with I elements using multiscale geometric analysis of the data such that every element in the data set may be represented, up to a certain precision ε , by m elements of the dictionary. The elements of the dictionary are called geometric wavelets. They share some similarities with the standard wavelets, but they are quite different from them in crucial ways. Geometric wavelets are not based on dilation and translation and their multiscale analysis is nonlinear. The authors develop fast but nonlinear algorithms for computing a fast geometric wavelet transform and discuss the problem of estimating the intrinsic dimension of a point cloud and give different examples to illustrate their techniques.

The last chapter that deals with data in high-dimensional spaces is Chap. 11 by Lieu and Saito. Many problems in pattern recognition often require comparison between ensembles of signals (i.e., points in a high-dimensional space) instead of comparing individual signals. Let

$$X^i = \{x_1^i, \dots, x_{m_i}^i\},$$

where x_j^i is a signal or a vector in \mathbb{R}^d . The set X^i is called a training ensemble, which is assumed to have a unique label among C possible labels, and m_i is the number

of signals in the ensemble X^i . The collection of ensembles $X = \cup_{i=1}^M X^i \subset \mathbb{R}^d$ is called a collection of M training ensembles. Let $\mathcal{Y} = \cup_{j=1}^N \mathcal{Y}^j \subset \mathbb{R}^d$ be a collection of test ensembles, where $\mathcal{Y}^j = \{Y_1^j, \dots, Y_{n_j}^j\}$ is the j th test ensemble of n_j signals. The goal is to classify each \mathcal{Y}^j to one of the possible C classes given the training ensembles X . The authors propose an algorithm for doing just that.

The proposed algorithm consists of two main steps. The first step performs the dimensionality reduction without losing important features of the training ensembles, followed by constructing a compact representation, called a signature, that represents an essence of the ensemble in the reduced space. The second step embeds a given test ensemble into the reduced space obtained in the first step followed by classifying it to the label of the training ensemble whose signature is most similar to that of the test ensemble. How to define the similarity or the distance measure between signatures is the key issue discussed in this work.

For the first step the authors try to find an effective low-dimensional representation of the data. The most well-known linear embedding techniques are Principal Component Analysis (PCA) and Multidimensional Scaling (MDS). More recently, many nonlinear methods for dimensionality reduction, such as Laplacian eigenmaps and diffusion maps, have been proposed in order to improve the shortcomings of PCA/MDS. The authors compare the performance of these reduction methods and then propose an algorithm that extends a given test ensemble into the trained embedding space. They then measure the distance between the test ensemble and each training ensemble in that space, and classify it using the nearest neighbor method. The label of the training ensemble whose signature is *closest* to that of a given test ensemble is assigned to it. This raises the question of how you define the distance measure for the nearest neighbor classifier. There are many such measures for comparing two given signatures, such as the usual Euclidean distance. But the most robust one the authors choose is the so-called Earth Mover's Distance (EMD).

The chapter is concluded by describing the results of numerical experiments on two real data sets to illustrate how the proposed algorithm can be applied in practice. The first experiment is on underwater object classification using acoustic signals and the second is on a lip-reading problem, whose objective is to train a machine to automatically recognize the spoken words from the movements of the lips captured on silent video segments (no sound is involved).

1.1.3 Wavelet Applications

The third part of the monograph is dedicated to wavelets applications. Some of the applications are based on the continuous wavelet transform and its generalizations. Recall that the continuous wavelet transform of a function $f \in L^2(\mathbb{R})$ is defined as

$$W_\psi[f](a, b) = \frac{1}{\sqrt{|a|}} \int_{\mathbb{R}} f(x) \psi\left(\frac{x-b}{a}\right) dx, \quad a, b \in \mathbb{R}, a \neq 0,$$

where ψ is the mother wavelet which is assumed to satisfy the admissibility condition

$$\int_{\mathbb{R}} \frac{|\hat{\psi}(\omega)|^2}{|\omega|} d\omega < \infty. \quad (1.4)$$

Spherical harmonics have been, for a long time, the main tool for analyzing functions on spheres but in recent years spherical wavelets have become a competing tool. Chapter 12 by D. Geller and A. Mayeli deals with the construction of wavelets on a smooth compact (connected) Riemannian manifold M , of dimension n , in particular, on spheres. They first show how to construct nearly tight frames that are well-localized both in frequency and space on M and then introduce what they call spin wavelets. They then apply their construction to the analysis of cosmic microwave background radiation (CMB).

Starting with an appropriately chosen mother wavelet on the real line and replacing the scale parameter in the associated continuous wavelet transform by a positive self-adjoint operator T on $L^2(M)$, an analogue of the wavelet transform can be defined. Since the transition from the real line to compact Riemannian manifolds is not very obvious, we will sketch the procedure.

From (1.4), we obtain for a suitable smooth function h with $h(0) = 0$

$$0 < \int_0^\infty \frac{|h(t)|^2}{t} dt = c < \infty,$$

which implies that for all $\xi > 0$

$$\int_0^\infty |h(t\xi)|^2 \frac{dt}{t} = c < \infty. \quad (1.5)$$

Standard discretization of this formula yields for $a > 1, \xi > 0$ (see [3, p. 69])

$$0 < A_a \leq \sum_{j \in \mathbb{Z}} |h(a^{2j}\xi)|^2 \leq B_a < \infty.$$

Now, let $h(\xi) = \xi^2 e^{-\xi^2}$, so that $\psi = \check{h}$ is a constant multiple of the Mexican hat function. By multiplying (1.5) by $|F(\xi)|^2$, where F is a square integrable function and integrating with respect to ξ , we have

$$\int_{\mathbb{R}} \|F * \psi_t\|_2^2 \frac{dt}{t} = c \|F\|_2^2,$$

where $\psi_t(x) = (1/t)\psi(x/t)$. The authors show that by discretizing the last integral, one obtains

$$A \|F\|_2^2 \leq \sum |\langle F, \phi_{j,k} \rangle|^2 \leq B \|F\|_2^2,$$

where

$$\phi_{j,k} = \sqrt{ba^j} \psi_{a^j}(x - bka^j),$$

for sufficiently small b and a sufficiently close to 1.

Replacing ξ in (1.5) by a positive self-adjoint operator T on a Hilbert space \mathcal{H} , the authors show that

$$\int_0^\infty |h|^2(tT) \frac{dt}{t} = c(I - P),$$

and

$$A_\alpha(I - P) \leq \sum_{j \in \mathbb{Z}} |h|^2(a^{2j}T) \leq B_\alpha(I - P),$$

where P is the projection onto the null space of T , I is the identity operator, A_α and B_α are constants.

As a special case, take $T = -d^2/dx^2$ on the real line, so that

$$[h(t^2T)F](x) = \int_{\mathbb{R}} K_t(x, y) F(y) dy, \quad F \in L^2,$$

where $K_t(x, y) = \psi_t(x - y)$. On a smooth compact (connected) Riemannian manifold M , T is taken as the Laplace–Beltrami operator. A wavelet theory can now be developed for $L^2(M)$. A characterization of the Besov spaces using the frames constructed earlier can be obtained. The authors proceed to introduce what they call spin wavelets on the sphere and discuss their applications.

One of the most interesting features of this work is its applications to physics, in particular, to cosmic microwave background radiation (CMB), which was emitted after the Big Bang and is regarded as one of the pieces of evidence for the Big Bang theory. In this application, spherical wavelets seem to work better than spherical harmonics since the former are well-localized.

P. Fisher and K. Tung in Chap. 13 present two applications of wavelets as a numerical tool for atmospheric data analysis. In the first application, they use the continuous wavelet transform to determine the local Quasi-Biennial Oscillation (QBO) period, which is a dominant oscillation of the equatorial stratospheric zonal wind. The period is irregular but averages to about 28 months. Because the continuous wavelet transform can determine the local intrinsic period of an oscillation, it gives a more objective method for calculating the period of the QBO than previously used subjective methods.

The second application employs a wavelet-based multifractal formalism to analyze multifractal signals and find their singularity spectra. The authors apply the wavelet-based multifractal approach to the analysis of two sets of atmospheric data.

In the last chapter of the monograph, N. Whitmal, J. Rutledge, and J. Cohen, present a wavelet based approach for denoising speech signals for digital hearing aids. They describe an algorithm to remove background noise from noisy speech signals for use in digital hearing aids. The algorithm, which is based on wavelets and wavelet packets, was the result of a series of successive experiments in which the algorithm was adjusted to overcome a number of problems, such as the selection of a basis that best distinguishes noise from speech and the choice of the optimal number of coefficients from which the denoised signal can be constructed. Experiments performed at a computer lab in the Northwestern University Department of

Electrical Engineering and Computer Science showed, via spectrograms and tables of signal to noise ratios, that their method provided better noise reduction than spectral subtraction or other subspace methods. Testing on hearing impaired subjects, however, yielded intelligibility results comparable with other methods. Subsequent research revealed, both theoretically and experimentally, that subspace methods that involve hard thresholding produce effects on normal hearing listeners that are similar to recruitment of loudness, a common hearing disorder.

References

1. X. Dai and D. Larson, *Wandering vectors for unitary systems and orthogonal wavelets*, Mem. Amer. Math. Soc. Vol. 134 (1998), No. 640, viii+68.
2. X. Dai, D. Larson, and D. Speegle, *Wavelet sets in \mathbb{R}^n* , J. Fourier Anal. Appl., Vol. 3 (1997), No. 4, pp. 451-456.
3. I. Daubechies, *Ten Lectures on Wavelets*, SIAM Publ. Philadelphia (1992).
4. I. Krishtal, B. Robinson, G. Weiss, and E. N. Wilson, *Some simple Haar-type wavelets in higher dimensions*, J. Geom. Anal., Vol. 17 (2007), pp. 87-96.
5. S. Mallat, *Multiresolution approximation and wavelets*, Trans. Amer. Math. Soc., Vol. 315 (1989), pp. 69-88.
6. Y. Meyer, *Ondelettes Et Opérateurs I: Ondelettes*, Hermann Éditeurs des Sciences et des Arts, Paris (1990).
7. J. O. Strömberg: *Wavelets in higher dimensions*. Documenta Mathematica Extra Volume ICM-1998(3), pp. 523-532.

Part I
The Mathematical Theory of Wavelets

Chapter 2

The Construction of Wavelet Sets

John J. Benedetto and Robert L. Benedetto

Abstract Sets Ω in d -dimensional Euclidean space are constructed with the property that the inverse Fourier transform of the characteristic function $\mathbf{1}_\Omega$ of the set Ω is a single dyadic orthonormal wavelet. The iterative construction is characterized by its generality, its computational implementation, and its simplicity. The construction is transported to the case of locally compact abelian groups G with compact open subgroups H . The best known example of such a group is $G = \mathbb{Q}_p$, the field of p -adic rational numbers (as a group under addition), which has the compact open subgroup $H = \mathbb{Z}_p$, the ring of p -adic integers. Fascinating intricacies arise. Classical wavelet theories, which require a non-trivial discrete subgroup for translations, do not apply to G , which may not have such a subgroup. However, our wavelet theory is formulated on G with new group theoretic operators, which can be thought of as analogues of Euclidean translations. As such, our theory for G is structurally cohesive and of significant generality. For perspective, the Haar and Shannon wavelets are naturally antipodal in the Euclidean setting, whereas their analogues for G are equivalent.

2.1 Introduction

2.1.1 Background

We shall give a general method for constructing single dyadic orthonormal wavelets, which generate wavelet orthonormal bases (ONBs) for the space L^2 of square-integrable functions in two important antipodal cases. The cases are $L^2(\mathbb{R}^d)$, where \mathbb{R}^d is d -dimensional Euclidean space, and $L^2(G)$, where G belongs to the class of locally compact Abelian groups (LCAGs) which contain a compact open subgroup, and which are often used in number theoretic applications. The method

J.J. Benedetto (✉)

Department of Mathematics, University of Maryland, College Park, MD 20742, USA

e-mail: jjb@math.umd.edu

and associated theory for $L^2(\mathbb{R}^d)$ were introduced by Manuel León, Songkiat Sumetkijakan, and the first named author in [20] (1999), [21] (2001), [22] (2002), [23] (2003). The theory for $L^2(G)$ was established by the authors in [17] (2004), [24] (2004).

This *constructive* method, which we refer to as the *neighborhood mapping construction* (NMC) was inspired by groundbreaking operator theoretic work due to Dai and Larson [34] (1998) and Dai, Larson, and Speegle [35] (1997). There was comparably compelling work, contemporaneous to [20] (1999), [21] (2001), in abstract harmonic analysis by Baggett, Medina, and Merrill [11] (1999).

The catalyst for our original research was a preprint of the Soardi–Weiland paper [91] (1998). The aforementioned, as well as less known but equally formidable results by Zakharov [98] (1996), were aimed at establishing the existence of *single dyadic orthonormal wavelets* ψ for $L^2(\mathbb{R}^d)$, $d > 1$, i.e., $\{\psi_{m,n} : m \in \mathbb{Z}, n \in \mathbb{Z}^d\}$ is an orthonormal basis (ONB) for $L^2(\mathbb{R}^d)$, where

$$\psi_{m,n}(x) = 2^{md/2} \psi(2^m x - n). \quad (2.1)$$

It turns out that the Fourier transform of such a function ψ is the characteristic function $\mathbf{1}_\Omega$ of a set Ω , and such sets and their generalizations are called *wavelet sets*. Besides describing the NMC, we shall give a significant list of references to illustrate a range of settings and problems associated with wavelet sets, and to provide perspective about the role and extent of the NMC in wavelet theory and its applications.

For some time there was doubt about the existence of single dyadic orthonormal wavelets ψ for \mathbb{R}^d , $d > 1$. In fact, the most common construction of wavelet ONBs was from the theory of multiresolution analysis (MRA) which requires $2^d - 1$ functions ψ_j , $j = 1, \dots, 2^d - 1$, to generate the resulting ONB, $\{(\psi_j)_{m,n}\}$, see [81] (1990), [38] (1992), [31] (1994), [76] (1994), [43] (1997), [78] (1998), [77] (1992), [80] (1986), [48] (1992), [95] (1994), [8] (1995), [47] (1995) for MRA theory on \mathbb{R}^d or, more generally on LCAGs containing cocompact discrete subgroups; cf. the work on minimally supported wavelets [42] (1996), [53] (1996), [54] (1997). Thus, the wavelets we construct, and those in [34] (1998), [35] (1997), [11] (1999) are not derived from any MRA. On the other hand, there are unifying general approaches, e.g., [85] (1998), [11] (1999), [12] (1999), [30] (1999), [84] (1993). There are also results on wavelet theory in a variety of natural settings such as Lie groups and manifolds, sometimes coupled with structural constraints such as MRA, e.g., [72] (1989), [32] (1995), [58] (1995), [97] (1996), [69] (1996), [53] (1996), [4] (1997), [60] (1997), [70] (1998), [71] (1998), [89] (1999), [7] (1999), [52] (1999), [3] (2000), [6] (2000), [65] (2002), [83] (2002), [5] (2004), [59] (2004), [66] (2004), [61] (2005), [45] (2005), [44] (2005), [90] (2009), [62] (2009), and classical work in harmonic analysis on local fields, e.g., [96] (1975) and [37] (1983). This list contains several papers dealing with the p -adics or other local fields on which we have focused in our wavelet analysis of $L^2(G)$, so we note the explicitness, generality, computability, and algebraic cohesiveness of our approach. This will be spelled out in Sects. 2.6–2.10.

Remark 2.1. a. Wavelet ONBs go far beyond the dyadic case. For example, the $d \times d$, dyadic diagonal matrix A (with 2 s along the diagonal), corresponding to (2.1), can be replaced by real expansive $d \times d$ matrices for which $A(\mathbb{Z}^d) \subseteq \mathbb{Z}^d$. As such, (2.1) can be replaced by functions of the form

$$(\psi_j^A)_{m,n}(x) = |\det(A)|^{m/2} \psi_j(A^m x - n), \quad (2.2)$$

where $j = 1, \dots, J$, $m \in \mathbb{Z}$, $n \in \mathbb{Z}^d$, e.g., see [80] (1986), [64] (1992), [25] (1999). We mention this, since we can define non-separable filters with corresponding matrix A and $J = 1$ to obtain a single MRA wavelet ψ^A for which $\{(\psi^A)_{m,n}\}$ is an ONB for $L^2(\mathbb{R}^d)$.

b. The reason we have not chosen this path to obtain single orthonormal wavelets, in spite of the elegance of MRA, is to make use of the “zooming” property of the dyadic case. In fact, by zooming-in and -out, because of powers of 2 (or of any $n \geq 2$), we can fathom multiscale phenomena in a function and/or control computational costs vis-à-vis signal resolution in reconstruction.

Remark 2.2. One aspect of the applicability alluded to in Remark 2.1.b is to provide another mathematical tool, along with dimension reduction techniques, e.g., with which to manage massive data sets generated by data-creating devices such as supercomputers, internet traffic, CAT scanners, and digital cameras. IDC estimates that the world generated 487 billion gigabytes of information in 2008. This creates formidable problems for obtaining digital representations suitable for storage, transmission, and/or recovery, as well as for handling information accurately, efficiently, and robustly. In the Epilogue, we comment on the process of useful implementation of single dyadic orthonormal wavelets for $L^2(\mathbb{R}^d)$, $d \gg 0$.

2.1.2 Notation and Outline

We shall employ the usual notation in harmonic analysis and wavelet theory as found in [15] (1997), [38] (1992), [82] (1992), and [94] (1971). The Fourier transform of the function $f: \mathbb{R}^d \rightarrow \mathbb{C}$ is formally defined by

$$\hat{f}(\gamma) = \int f(x) e^{-2\pi i x \cdot \gamma} dx,$$

where \int denotes integration over \mathbb{R}^d ; and the inverse Fourier transform F^\vee of $F: \hat{\mathbb{R}}^d \rightarrow \mathbb{C}$ is formally defined by

$$F^\vee(x) = \int F(\gamma) e^{2\pi i x \cdot \gamma} d\gamma, \quad x \in \mathbb{R}^d,$$

where $\hat{\mathbb{R}}^d$ is \mathbb{R}^d considered as the spectral domain. \mathbb{Z} is the ring of integers and \mathbb{T} designates the quotient group $\hat{\mathbb{R}}/\mathbb{Z}$. If F is a 1-periodic function on $\hat{\mathbb{R}}$ with Fourier series $S(F)(\gamma) = \sum c_n e^{-2\pi i n \gamma}$, then the Fourier coefficients c_n are designated

by $F^\vee[n]$. Further, translation of a function f by x is designated by $\tau_x f$, i.e., $\tau_x f(y) = f(x - y)$. Finally, if $\Omega \subseteq \widehat{\mathbb{R}}^d$, then its Lebesgue measure is denoted by $|\Omega|$. The term *measurable* will mean *Lebesgue measurable*.

The paper is structured as follows. Sections 2.2–2.5 deal with the Euclidean theory of wavelet sets and Sects. 2.6–2.10 deal with the non-Euclidean theory. Section 2.11, the Epilogue, briefly broadens some of the conventional perspective about wavelet sets and their genuine applicability. Generally, we refer to our original papers for the proofs of theorems. However, there are a few salient exceptions related to our opinion of what constitutes general interest, or where we deem the details or structure of the proof to be particularly informative or surprising. In addition, we present many examples.

Section 2.2 is devoted to the geometry of Euclidean wavelet sets, as well as to fundamental roots based in Lusin’s conjecture (and thus Carleson’s theorem) and the Littlewood–Paley theory. Section 2.3 provides the details of our neighborhood mapping construction (NMC) of wavelet sets. It is highly motivated geometrically, but ultimately rather intricate. In Sect. 2.4, we prove a basic theorem about frame wavelet sets which we view as a major means of applying wavelet sets in a host of signal processing applications dealing with large data sets. Finally, in Sect. 2.5, for the Euclidean theory, we give geometrical examples with suggestive topological implications, as well as structural implications of the NMC and a hint of the breadth and beauty of NMC constructible wavelet sets.

Earlier we were intrigued by the possibility and utility of wavelet sets in number theory, based on one of the author’s ideas about idelic pseudo-measures [14] (1979), [13] (1973). Sections 2.6–2.10 are our foray into this area. The background dealing with LCAGs, the p -adic field, and generic wavelet theory in this setting is the subject of Sect. 2.6. Our fundamental idea to ensure the mathematical cohesiveness and resulting group theoretic canonicity and mathematical beauty of our approach is the subject of Sect. 2.7. With this background, Sect. 2.8 gives a basic geometrical result for the number theoretic setting analogous to the point of view of Sect. 2.2. This substantive theory is the background for the number theoretic construction and algorithm of Sect. 2.9, which itself is driven by the ideas of Sect. 2.3. Finally, in Sect. 2.10, we give examples indicating the incredible breadth of the number theoretic NMC.

2.2 Geometry of Euclidean Wavelet Sets

2.2.1 Wavelet Sets, Tilings, and Congruences

A set Ω , together with the property that $\psi = \mathbf{1}_\Omega^\vee$ is a single dyadic orthonormal wavelet, is a *wavelet set*. Our construction of such sets Ω is the subject of Sect. 2.3, and our basic geometrical approach is not unrelated to constructions of Leonardo da Vinci and Maurits C. Escher.

Remark 2.3. Consider $\Psi = \{\psi^1, \psi^2, \dots, \psi^M\} \subseteq L^2(\mathbb{R}^d)$. We say Ψ is a *set of wavelet generators* for $L^2(\mathbb{R}^d)$ if

$$\{\psi_{m,n}^i(\cdot) = 2^{md/2} \psi^i(2^m \cdot -n) : m \in \mathbb{Z}, n \in \mathbb{Z}^d, i = 1, \dots, M\}$$

is an ONB for $L^2(\mathbb{R}^d)$. Auscher [8] (1995) proved that every set of wavelet generators for $L^2(\mathbb{R})$, whose members satisfy a weak smoothness and decay condition on the Fourier transform side, must come from an MRA. Further, it is known, e.g., see [8] (1995), [82] (1992), that for a given dyadic MRA there is a wavelet collection consisting of $2^d - 1$ elements. There is an analogous assertion for the expansive matrix case.

Because of this remark, and notwithstanding Journé's celebrated example of a non-MRA wavelet basis for $L^2(\mathbb{R})$, e.g., [38] (1992), there was some question during the mid-1990s about the existence of multidimensional single dyadic orthonormal wavelets. Dai, Larson, and Speegle [35] (1997), referenced earlier, proved the existence of such wavelets in $L^2(\mathbb{R}^d)$, $d > 1$. Their proof depended on wavelet sets and used operator algebra methods. Some of the initial reaction was a combination of disbelief and disinterest, the latter response due to the prevailing intuition that such wavelets would be difficult to implement in an effective way.

- Definition 2.1.** *a.* Let $\Omega \subseteq \widehat{\mathbb{R}}^d$ be measurable. A *tiling* of Ω is a collection $\{\Omega_l : l \in \mathbb{Z}\}$ of measurable subsets of $\widehat{\mathbb{R}}^d$ such that $\bigcup_l \Omega_l$ and Ω differ by a set of measure 0, and, for all $l \neq j$, $|\Omega_l \cap \Omega_j| = 0$.
- b.* Let $\Omega, \Theta \subseteq \widehat{\mathbb{R}}^d$ be measurable. If there exist a tiling $\{\Omega_l : l \in \mathbb{Z}\}$ of Ω and a sequence $\{k_l : l \in \mathbb{Z}\} \subseteq \mathbb{Z}^d$ such that $\{\Omega_l + k_l : l \in \mathbb{Z}\}$ is a tiling of Θ , then Ω and Θ are \mathbb{Z}^d -translation congruent or τ -congruent. This is equivalent to the existence of tilings $\{\Omega_l : l \in \mathbb{Z}\}$ and $\{\Theta_l : l \in \mathbb{Z}\}$ of Ω and Θ , respectively, and a sequence $\{n_l : l \in \mathbb{Z}\} \subseteq \mathbb{Z}^d$ such that $\Omega_l = \Theta_l + n_l$, for all $l \in \mathbb{Z}$.
- c.* Let $\Omega, \Theta \subseteq \widehat{\mathbb{R}}^d$ be measurable. If there exist a tiling $\{\Omega_l : l \in \mathbb{Z}\}$ of Ω and a sequence $\{m_l : l \in \mathbb{Z}\} \subseteq \mathbb{Z}$, where $\{2^{m_l} \Omega_l : l \in \mathbb{Z}\}$ is a tiling of Θ , then Ω and Θ are dyadic-dilation congruent or δ -congruent. This is equivalent to the existence of tilings $\{\Omega_l : l \in \mathbb{Z}\}$ and $\{\Theta_l : l \in \mathbb{Z}\}$ of Ω and Θ , respectively, and a sequence $\{m_l : l \in \mathbb{Z}\} \subseteq \mathbb{Z}$ such that $\Omega_l = 2^{m_l} \Theta_l$, for all $l \in \mathbb{Z}$.
- d.* We shall deal with tilings of $\widehat{\mathbb{R}}^d$ by translation or dilation of a measurable set $\Omega \subseteq \widehat{\mathbb{R}}^d$. Thus, $\{\Omega + n : n \in \mathbb{Z}^d\}$ is a *tiling* of $\widehat{\mathbb{R}}^d$ means that $|\widehat{\mathbb{R}}^d \setminus \bigcup_{n \in \mathbb{Z}^d} (\Omega + n)| = 0$ and $|(\Omega + m) \cap (\Omega + n)| = 0$ when $m \neq n$. Similarly, $\{2^m \Omega : m \in \mathbb{Z}\}$ is a *tiling* of $\widehat{\mathbb{R}}^d$ means that $|\widehat{\mathbb{R}}^d \setminus \bigcup_{m \in \mathbb{Z}} (2^m \Omega)| = 0$ and $|(2^j \Omega) \cap (2^m \Omega)| = 0$ when $j \neq m$.
- e.* It is not difficult to see that the concept of Ω being \mathbb{Z}^d -translation congruent to $[-\frac{1}{2}, \frac{1}{2}]^d$ is equivalent to $\{\Omega + n : n \in \mathbb{Z}^d\}$ being a tiling of $\widehat{\mathbb{R}}^d$.

Remark 2.4. *a.* The notion of congruence plays a role in several facets of wavelet theory besides the results in this paper. Congruence criteria were used by Albert Cohen in 1990 to characterize the orthonormality of scaling functions defined by infinite products of dilations of a quadrature mirror filter, e.g., [38, pp.182–186]

(1992). The same notion of congruence also plays a fundamental role in work on self-similar tilings of \mathbb{R}^d by Gröchenig, Haas, Lagarias, Madych, Yang Wang, et al., e.g., [67] (1997), [68] (2000).

- b. The notion of \mathbb{Z}^d -translation congruence is intrinsically related to bijective restrictions of the canonical surjection $h : G \longrightarrow G/H$, where G is a locally compact group and H is a closed subgroup. An analysis of this relation is found in [16, Sect. 3] (1998) in the context of Kluvánek's sampling theorem for locally compact Abelian groups. Kluvánek's sampling formula for a signal f quantitatively relates the sampling rate with the measure of the subsets of a given bandwidth corresponding to the frequency content of f .

Wavelet sets and tilings are related by the following theorem. For an elementary proof, as well as a more complicated one, see [20] (1999). The existence of wavelet sets is not obvious, and this is the point of Sect. 2.3.

Theorem 2.1. *Let $\Omega \subseteq \widehat{\mathbb{R}}^d$ be a measurable set. Ω is a wavelet set if and only if*

- i. $\{\Omega + n : n \in \mathbb{Z}^d\}$ is a tiling of $\widehat{\mathbb{R}}^d$, and
- ii. $\{2^m \Omega : m \in \mathbb{Z}\}$ is a tiling of $\widehat{\mathbb{R}}^d$.

Corollary 2.1. *Let $\Omega \subseteq \widehat{\mathbb{R}}^d$. Ω is a wavelet set if and only if Ω is \mathbb{Z}^d -translation congruent to $[0, 1)^d$ and Ω is dyadic-dilation congruent to $[-1, 1)^d \setminus [-\frac{1}{2}, \frac{1}{2})^d$.*

Definition 2.2. A collection $\Omega^1, \dots, \Omega^L$ of measurable subsets of $\widehat{\mathbb{R}}^d$ is a *wavelet collection of sets* if $\{\mathbf{1}_{\Omega^1}^\vee, \dots, \mathbf{1}_{\Omega^L}^\vee\}$ is a set of wavelet generators for $L^2(\mathbb{R}^d)$.

We have the following generalization of Theorem 2.1. It should be compared with Theorem 2.16, whose more complicated proof is included.

Theorem 2.2. *Let $\Omega^1, \dots, \Omega^L$ be pairwise disjoint measurable subsets of $\widehat{\mathbb{R}}^d$. The family $\{\Omega^l : l = 1, \dots, L\}$ is a wavelet collection of sets if and only if each $|\Omega^l| = 1$ and the following conditions are satisfied:*

- i. *For each fixed $l = 1, \dots, L$, $\{\Omega^l + k : k \in \mathbb{Z}^d\}$ is a tiling of $\widehat{\mathbb{R}}^d$;*
- ii. *If $\Omega = \bigcup_{l=1}^L \Omega^l$, then $\{2^j \Omega : j \in \mathbb{Z}\}$ is a tiling of $\widehat{\mathbb{R}}^d$.*

Remark 2.5. In light of our dyadic results in this paper involving functions of the form $\mathbf{1}_\Omega^\vee$, we point out that Gu and Han [49] (2000) proved that, in the setting of (2.2), if $|\det A| = 2$, then there is a measurable set $\Omega \subseteq \widehat{\mathbb{R}}^d$ such that $\{2^{md/2} \mathbf{1}_\Omega^\vee(A^m x - n) : m \in \mathbb{Z}, n \in \mathbb{Z}^d\}$ is an ONB for $L^2(\mathbb{R}^d)$. This result can be viewed as a converse of the following theorem: if $\psi \in L^2(\mathbb{R}^d)$ is a single wavelet constructed from an MRA associated with (\mathbb{Z}^d, A) , then $|\det A| = 2$, see [9] (1995) and [50] (1997), cf. [64] (1992) and [28] (1993).

2.2.2 Kolmogorov Theorem and Littlewood–Paley Wavelet ONB

In 1922, Kolmogorov [63] (1924) proved that if $F \in L^2(\mathbb{T})$ and $S_N(F)$ is the N th partial sum of the Fourier series $S(F)$ of F , then

$$\lim_{n \rightarrow \infty} S_{2^n}(F)(\gamma) = F(\gamma) \text{ a.e.} \quad (2.3)$$

His proof is elementary, short, and clever; and the result is still valid when $\{2^n\}$ is replaced by more general lacunary sequences. Writing

$$\Delta_j F(\gamma) = \sum_{2^j \leq |n| < 2^{j+1}} F^\vee[n] e^{-2\pi i n \gamma}, \quad j = 0, 1, \dots,$$

Equation (2.3) can be restated as

$$F(\gamma) = F^\vee[0] + \sum_{j=0}^{\infty} \Delta_j F(\gamma) \text{ a.e.}, \quad (2.4)$$

which can be interpreted as a frequency decomposition of F based on dyadic intervals. Equation (2.4) plays a fundamental role in Littlewood–Paley theory, and was stated and proved as a consequence of deep results in their theory in the setting of $L^p(\mathbb{T})$, $p > 1$, see [74, Theorem 5] (1931) and [75, Theorem 8] (1937). The Littlewood–Paley theory is an important part of twentieth century harmonic analysis, e.g., see [29] (1978), [41] (1977), [46] (1991), [92, Chap. 14] (1970), and [93] (1970).

From our point of view, (2.4) can be adjusted to incorporate time-frequency localization, at least within the constraints of the classical uncertainty principle; and it can be thought of as a primordial wavelet decomposition, e.g., [82, pp.19–20] (1992). In fact, in the setting of \mathbb{R} , the decomposition (2.4), properly localized in time and reformulated in terms of multiresolution analysis, becomes the *Littlewood–Paley* or *Shannon* wavelet orthonormal basis decomposition

$$f = \sum_{m,n} \langle f, \psi_{m,n} \rangle \psi_{m,n}, \quad \text{for all } f \in L^2(\mathbb{R}), \quad (2.5)$$

where

$$\hat{\psi} = \mathbf{1}_\Omega, \quad \Omega = \left[-1, -\frac{1}{2}\right) \cup \left[\frac{1}{2}, 1\right)$$

is the Fourier transform of the Littlewood–Paley or Shannon wavelets ψ .

The decomposition (2.5) can be proved in several standard ways, but the most convenient is to combine the orthonormality of $\{\psi_{m,n}\}$ with the fact that

$$\sum_{m,n} |\langle f, \psi_{m,n} \rangle|^2 = \|f\|_{L^2(\mathbb{R})}^2, \quad \text{for all } f \in L^2(\mathbb{R}), \quad (2.6)$$

e.g., see [38, pp. 115-16] (1992) and [46] (1991) for further details. The proof of Equation (2.6) is the calculation

$$\begin{aligned}
\sum |\langle f, \psi_{m,n} \rangle|^2 &= \sum_{m,n} 2^m \left| \int_{\Omega} \hat{f}(2^m \lambda) e^{2\pi i n \lambda} d\lambda \right|^2 \\
&= \sum_{m,n} \left| \int_{-\frac{1}{2}}^{\frac{1}{2}} \left(\hat{f}(2^m(\gamma-1)) \mathbf{1}_{[-1, -\frac{1}{2})}(\gamma-1) \right. \right. \\
&\quad \left. \left. + \hat{f}(2^m(\gamma+1)) \mathbf{1}_{[\frac{1}{2}, 1)}(\gamma+1) \right) e^{2\pi i n \gamma} d\gamma \right|^2 \\
&= \sum_m \int_{-\frac{1}{2}}^{\frac{1}{2}} \left| \hat{f}(2^m(\gamma-1)) \mathbf{1}_{[-1, -\frac{1}{2})}(\gamma-1) \right. \\
&\quad \left. + \hat{f}(2^m(\gamma+1)) \mathbf{1}_{[\frac{1}{2}, 1)}(\gamma+1) \right|^2 d\gamma \\
&= \sum_m 2^m \int_{\Omega} |\hat{f}(2^m \lambda)|^2 d\lambda = \sum_m \int_{2^m \Omega} |\hat{f}(\gamma)|^2 d\gamma = \|f\|_{L^2(\mathbb{R})}^2. \quad (2.7)
\end{aligned}$$

The points to be made are that (2.7) is essentially a geometrical argument, and also that it can be generalized. The fact that (2.7) is a geometrical argument is immediate from the second equality, which depends on the \mathbb{Z} -translation congruence of Ω and $[-\frac{1}{2}, \frac{1}{2}]$, and the last equality, which is due to the fact that $\{2^m \Omega\}$ is a tiling of $\widehat{\mathbb{R}}$.

Thus, the Shannon wavelet ψ does in fact give rise to a dyadic wavelet ONB for $L^2(\mathbb{R})$. Moreover, $\widehat{\psi} = \mathbf{1}_{\Omega}$ so we are dealing with the wavelet set $\Omega = [-1, -\frac{1}{2}) \cup [\frac{1}{2}, 1)$; and, most important, the proof that $\{\psi_{m,n}\}$ is an ONB for $L^2(\mathbb{R})$ depends entirely on the tiling criteria of Theorem 2.1.

In the next section, we shall give a general construction of wavelet sets motivated by the tiling criteria of Theorem 2.1. Intuitively, these criteria assert that Ω must have fundamental characteristics of both squares *and* annuli.

2.3 The Construction of Euclidean Wavelet Sets

2.3.1 The Basic Construction

Let $\Omega_0 \subseteq [-N, N]^d \subseteq \widehat{\mathbb{R}}^d$ be a neighborhood of the origin with Lebesgue measure $|\Omega_0| = 1$, and further assume that Ω_0 is \mathbb{Z}^d -translation congruent to $[-\frac{1}{2}, \frac{1}{2}]^d$. Ω_0 will be iteratively transformed by the action of a mapping

$$T : \Omega_0 \longrightarrow \overline{[-2N, 2N]^d} \setminus [-N, N]^d$$

for some fixed N , where T is defined by the property that, for each fixed $\gamma \in \Omega_0$, $T(\gamma) = \gamma + k_{\gamma}$ for some $k_{\gamma} \in \mathbb{Z}^d$.

Because of the requirements of our forthcoming construction, we shall assume that the mapping T , defined in terms of the translation property $T\gamma = \gamma + k_\gamma$, also has the properties that it is a measurable, injective mapping on Ω_0 , see [21, Proposition 3.1] (2001).

Algorithm 2.1. We now describe our original NMC construction of wavelet sets Ω depending on Ω_0 , N , and T . Let

$$\Lambda_0 = \Omega_0 \cap \left(\bigcup_{j \geq 1} 2^{-j} \Omega_0 \right) \text{ and } \Omega_1 = (\Omega_0 \setminus \Lambda_0) \cup T\Lambda_0.$$

Then,

$$\Omega_0 \setminus \Lambda_0 \subseteq \Omega_0 \text{ and } T\Lambda_0 \subseteq \overline{[-2N, 2N]^d \setminus [-N, N]^d}.$$

Next, let

$$\Lambda_1 = \Omega_1 \cap \left(\bigcup_{j \geq 1} 2^{-j} \Omega_1 \right),$$

and let $\Omega_2 = ((\Omega_0 \setminus \Lambda_0) \setminus \Lambda_1) \cup T\Lambda_0 \cup T\Lambda_1$. Then,

$$(\Omega_0 \setminus \Lambda_0) \setminus \Lambda_1 \subseteq \Omega_0$$

and

$$T\Lambda_0 \cup T\Lambda_1 \subseteq \overline{[-2N, 2N]^d \setminus [-N, N]^d}.$$

Notationally, we set $((\Omega_0 \setminus \Lambda_0) \setminus \Lambda_1) = \Omega_0 \setminus \Lambda_0 \setminus \Lambda_1$. Generally, for a given Ω_n , let

$$\Lambda_n = \Omega_n \cap \left(\bigcup_{j \geq 1} 2^{-j} \Omega_n \right),$$

and set

$$\Omega_{n+1} = (\Omega_0 \setminus \Lambda_0 \setminus \Lambda_1 \setminus \cdots \setminus \Lambda_n) \cup (T\Lambda_0 \cup T\Lambda_1 \cup \cdots \cup T\Lambda_n). \quad (2.8)$$

Then,

$$\Omega_0 \setminus \Lambda_0 \setminus \Lambda_1 \setminus \cdots \setminus \Lambda_n \subseteq \Omega_0$$

and

$$T\Lambda_0 \cup T\Lambda_1 \cup \cdots \cup T\Lambda_n \subseteq \overline{[-2N, 2N]^d \setminus [-N, N]^d}.$$

We define Ω as

$$\Omega = \left(\Omega_0 \setminus \bigcup_{k=0}^{\infty} \Lambda_k \right) \cup \left(\bigcup_{n=0}^{\infty} T\Lambda_n \right). \quad (2.9)$$

Denoting

$$\Omega_n^- = \Omega_0 \setminus \Lambda_0 \setminus \Lambda_1 \setminus \cdots \setminus \Lambda_{n-1}$$

and

$$\Omega_n^+ = T\Lambda_0 \cup T\Lambda_1 \cup \cdots \cup T\Lambda_{n-1},$$

we have

$$\Omega_n = \Omega_n^- \cup \Omega_n^+, \quad \Omega = \left(\bigcap_{n=0}^{\infty} \Omega_n^- \right) \cup \left(\bigcup_{n=0}^{\infty} \Omega_n^+ \right), \text{ and } |\Omega_n| = |\Omega_n^-| + |\Omega_n^+| = 1.$$

Thus, the set Ω is obtained by removing from Ω_0 all the Λ_i s and sending these disjoint subsets into $\overline{[-2N, 2N]^d \setminus [-N, N]^d}$ by means of the mapping T . It should be noted that Ω is \mathbb{Z}^d -translation congruent to Ω_0 .

Theorem 2.3. Ω defined by (2.9) is a wavelet set, see [20] (1999), [21] (2001).

The following is the generalization of Theorem 2.3 corresponding to the geometrical characterization of Theorem 2.2.

Theorem 2.4. Let $\{\Omega_0^1, \dots, \Omega_0^L\} \subseteq \widehat{\mathbb{R}}^d$, and assume T and each Ω_0^l satisfy the hypotheses of Algorithm 2.1. Let $\{\Omega^1, \dots, \Omega^L\}$ be the sequence of sets constructed in Algorithm 2.1. Then, $\{\Omega^1, \dots, \Omega^L\}$ is a wavelet collection of sets, i.e.,

$$\{\psi^l : \mathbf{1}_{\Omega^l}^\vee : l = 1, \dots, L\}$$

is a set of wavelet generators for $L^2(\mathbb{R}^d)$.

2.3.2 A Generalization of the Neighborhood-Mapping Construction

It is assumed in the original NMC of Sect. 2.3.1 that Ω_0 is contained in $[-N, N]^d$ and that the range of the mapping T is contained in $\overline{[-2N, 2N]^d \setminus [-N, N]^d}$. As it turns out, this assumption on the range of T is not necessary. The purpose of the mapping T should only be to move the sets Λ_n , defined below, out of Ω_0 . In this section we prove that the procedure produces wavelet sets for a more general class of mappings T , thereby obtaining wavelet sets that we had not been able to obtain by the original construction.

Let Ω_0 be a bounded neighborhood of the origin that is \mathbb{Z}^d -translation congruent to the unit cube $Q = [-\frac{1}{2}, \frac{1}{2}]^d$. We shall consider measurable mappings $T : \widehat{\mathbb{R}}^d \rightarrow \widehat{\mathbb{R}}^d$ satisfying the following properties.

i. T is a \mathbb{Z}^d -translated mapping, i.e.,

$$\forall \gamma \in \widehat{\mathbb{R}}^d, \exists n_\gamma \in \mathbb{Z}^d \text{ such that } T(\gamma) = \gamma + n_\gamma.$$

ii. T is injective.

- iii. The range of $T - I$ is bounded, where I is the identity mapping on $\widehat{\mathbb{R}}^d$.
 iv. $[\bigcup_{i=1}^{\infty} T^i \Omega_0] \cap [\bigcup_{j=0}^{\infty} 2^{-j} \Omega_0] = \emptyset$, where $T^0 = I$ and $T^i = \underbrace{T \circ \cdots \circ T}_{i\text{-fold}}$.

Compared to the original NMC, the first two conditions on T are unchanged, while the last two relax the earlier assumption on the range of T . Condition *iii* says that $T(\gamma) = \gamma + n_\gamma$ cannot be arbitrarily far from γ . There must be a uniform bound on how far γ moves to $T(\gamma)$ but the range of T does not necessarily lie inside some square box. What condition *iv* says is that for any $\gamma \in \Omega_0$ the sequence $\{T(\gamma), T^2(\gamma), \dots, T^n(\gamma), \dots\}$ never returns to Ω_0 or any $2^{-j}\Omega_0, j > 1$. This weakens the earlier artificial assumption that T has to move points in Ω_0 out of a square containing Ω_0 .

Algorithm 2.2. Let T satisfy conditions *i*–*iv*. According to [21] (2001), [22] (2002), we iteratively construct a sequence of sets Ω_n each of which is \mathbb{Z}^d -translation congruent to Ω_0 , and hence tiles $\widehat{\mathbb{R}}^d$ by \mathbb{Z}^d -translates, as follows. For each $n = 0, 1, \dots$, we define

$$\Lambda_n = \Omega_n \cap \left[\bigcup_{j=1}^{\infty} 2^{-j} \Omega_n \right]$$

and

$$\Omega_{n+1} = (\Omega_n \setminus \Lambda_n) \cup T\Lambda_n. \quad (2.10)$$

This set Ω_{n+1} , defined by (2.10), is the same as the set Ω_{n+1} , defined by (2.8). However, by property *iv* and some set theoretic implications of it, we calculate that

$$\Omega_{n+1} = \left[\Omega_0 \setminus \bigcup_{i=0}^n \Lambda_i \right] \cup \left[\bigcup_{k=0}^n \left(T\Lambda_k \setminus \bigcup_{i=k+1}^n \Lambda_i \right) \right]. \quad (2.11)$$

Because of (2.11), we define the set

$$\Omega = \left[\Omega_0 \setminus \bigcup_{i=0}^{\infty} \Lambda_i \right] \cup \left[\bigcup_{k=0}^{\infty} \left(T\Lambda_k \setminus \bigcup_{i=k+1}^{\infty} \Lambda_i \right) \right]. \quad (2.12)$$

Theorem 2.5. Ω defined by (2.12) is a wavelet set.

The proof is found in [23] (2006). It depends on several useful implications of properties *i*–*iv* including the following result.

Proposition 2.1. For each $n = 0, 1, \dots$, we have the following.

- Ω_n tiles $\widehat{\mathbb{R}}^d$ by \mathbb{Z}^d -dyadic translation, and
- $\Omega_n \setminus \Lambda_n$ tiles $\widehat{\mathbb{R}}^d$ by dyadic-dilation.

2.4 Frame Wavelets

The concept of frame introduced by Duffin and Schaeffer [40] (1952) is a natural generalization of an orthonormal basis.

Definition 2.3. *a.* A countable family $\{\phi_i\}_{i \in \mathbb{Z}}$ of functions in $L^2(\mathbb{R}^d)$ is a *frame* for $L^2(\mathbb{R}^d)$ if there exist constants $0 < A \leq B < \infty$ called *frame bounds* for which

$$A\|f\|^2 \leq \sum_{i \in \mathbb{Z}} |\langle f, \phi_i \rangle|^2 \leq B\|f\|^2, \quad \text{for all } f \in L^2(\mathbb{R}^d).$$

When the frame bounds A and B coincide, $\{\phi_i\}$ is called a *tight frame*. If $A = B = 1$, the inequality becomes the Parseval identity and $\{\phi_i\}$ is aptly called a *Parseval frame*, i.e.,

$$\sum_{i \in \mathbb{Z}} |\langle f, \phi_i \rangle|^2 = \|f\|^2, \quad \text{for all } f \in L^2(\mathbb{R}^d).$$

- b.* An $L^2(\mathbb{R}^d)$ function ψ is a *frame wavelet*, respectively, *tight frame wavelet* and *Parseval frame wavelet*, if the generated family $\{\psi_{m,n} : m \in \mathbb{Z}, n \in \mathbb{Z}^d\}$ is a frame, respectively, tight frame and Parseval frame.
- c.* $\Omega \subseteq \widehat{\mathbb{R}}^d$ is a *(Parseval) frame wavelet set* if $\psi = \mathbf{1}_\Omega^\vee$ is a (Parseval) frame wavelet.

The following is the analogue of Theorem 2.1 for the case of Parseval frame wavelets.

Theorem 2.6. *Let $\Omega \subseteq \widehat{\mathbb{R}}^d$ be measurable. The following are equivalent.*

- i.* Ω is a Parseval frame wavelet set.
- ii.* Ω is \mathbb{Z}^d -translation congruent to a subset of $[0, 1)^d$ and Ω is dyadic-dilation congruent to $[-1, 1)^d \setminus [-\frac{1}{2}, \frac{1}{2})^d$.
- iii.* $\{\Omega + k : k \in \mathbb{Z}^d\}$ is a tiling of a subset of $\widehat{\mathbb{R}}^d$ and $\{2^n \Omega : n \in \mathbb{Z}\}$ is a tiling of $\widehat{\mathbb{R}}^d$.

In recent years, frame wavelets, in particular tight and Parseval frame wavelets, have been studied extensively, e.g., see [50] (1997) by Bin Han, as well as a related paper by Dai, Diao, Gu, and Deguang Han, [33] (2002).

It is a natural question to ask whether the sets Ω_n constructed from finite iterations in the NMC of Sect. 2.3 give rise to frame wavelets, respectively, tight frame wavelets and Parseval frame wavelets, i.e., whether the functions $\mathbf{1}_{\Omega_n}^\vee$ are frame wavelets, respectively, tight frame wavelets and Parseval frame wavelets. It turns out that $\mathbf{1}_{\Omega_n}^\vee$ is a frame wavelet with frame bounds 1 and 2, while we obtain Parseval frame wavelets from the auxiliary sets $\Omega_n \setminus \Lambda_n$ (Theorem 2.9 below). We shall prove this result not only because it is a bit surprising but also because it facilitates the genuine implementation of wavelet set theory.

To this end, we need the characterization of tight frame wavelets due to Bin Han [50] (1997) as well as to Ron and Shen (1997), Bownik (2000), Chui and Shi (2000), and Chui et al. [27] (2002) (Theorem 2.7 below). We also require Daubechies' sufficient condition for a function to be a frame wavelet [38] (1992), combined with

refinements by Kugarajah and Zhang (1995) and Hernández and Weiss [55] (1996) (Theorem 2.8 below). The proof for \mathbb{R} of Theorem 5.1 in [26] (2001) can be easily generalized to \mathbb{R}^d , and this generalization is Theorem 2.8.

Theorem 2.7. *Let $\psi \in L^2(\mathbb{R}^d)$. The family*

$$\{\psi_{m,n} = 2^{md/2} \psi(2^m \cdot -n) : m \in \mathbb{Z}, n \in \mathbb{Z}^d\}$$

is a Parseval frame if and only if

$$\sum_{j \in \mathbb{Z}} |\widehat{\psi}(2^j \xi)|^2 = 1 \quad \text{and} \quad t_q(\xi) = \sum_{j=0}^{\infty} \widehat{\psi}(2^j \xi) \overline{\widehat{\psi}(2^j(\xi + q))} = 0$$

for almost every $\xi \in \widehat{\mathbb{R}^d}$ and for all $q \in \mathbb{Z}^d \setminus 2\mathbb{Z}^d$.

Theorem 2.8. *Let $a > 1$, $b > 0$, and $\psi \in L^2(\mathbb{R}^d)$ be given. Suppose that*

$$A = \inf_{\|\xi\| \in [1, a]} \left[\sum_{n \in \mathbb{Z}} |\widehat{\psi}(a^n \xi)|^2 - \sum_{k \neq 0} \sum_{n \in \mathbb{Z}} |\widehat{\psi}(a^n \xi) \widehat{\psi}(a^n \xi + k/b)| \right] > 0,$$

$$B = \sup_{\|\xi\| \in [1, a]} \left[\sum_{n \in \mathbb{Z}} |\widehat{\psi}(a^n \xi)|^2 + \sum_{k \neq 0} \sum_{n \in \mathbb{Z}} |\widehat{\psi}(a^n \xi) \widehat{\psi}(a^n \xi + k/b)| \right] < \infty.$$

Then $\{a^{jd/2} \psi(a^j \cdot -kb)\}_{j \in \mathbb{Z}, k \in \mathbb{Z}^d}$ is a frame for $L^2(\mathbb{R}^d)$ with frame bounds A/b^d , B/b^d .

We begin the proof of Theorem 2.9 with the following lemma which uses Theorem 2.7. The 1-dimensional version of Lemma 2.1 first appeared in Theorem 4.1 of [50] (1997).

Lemma 2.1. *If a measurable set Θ tiles $\widehat{\mathbb{R}^d}$ by dyadic-dilation and $\Theta \subseteq \Omega$ for some measurable set Ω that tiles $\widehat{\mathbb{R}^d}$ by \mathbb{Z}^d -translation, then the function $\psi \in L^2(\mathbb{R}^d)$ defined by $\widehat{\psi} = \mathbf{1}_\Theta$ is a Parseval frame wavelet.*

Proof. Since Θ tiles $\widehat{\mathbb{R}^d}$ by dyadic dilation a.e., we have

$$\sum_{j \in \mathbb{Z}} |\widehat{\psi}(2^j \xi)|^2 = \sum_{j \in \mathbb{Z}} \mathbf{1}_\Theta(2^j \xi) = \mathbf{1}_{\bigcup_{j \in \mathbb{Z}} 2^{-j} \Theta}(\xi) = 1 \quad \text{a.e.}$$

We then compute

$$\begin{aligned} t_q(\xi) &= \sum_{j=0}^{\infty} \widehat{\psi}(2^j \xi) \overline{\widehat{\psi}(2^j(\xi + q))} = \sum_{j=0}^{\infty} \mathbf{1}_{2^{-j} \Theta}(\xi) \mathbf{1}_{2^{-j} \Theta - q}(\xi) \\ &= \sum_{j=0}^{\infty} \mathbf{1}_{2^{-j} [\Theta \cap (\Theta - 2^j q)]}(\xi). \end{aligned}$$

From the second assumption, $\Theta \cap (\Theta - 2^j q) \subseteq \Omega \cap (\Omega - 2^j q) = \emptyset$. Therefore, $t_q = 0$. Hence, by Theorem 2.7, ψ is a Parseval frame wavelet. \square

Theorem 2.9. *For each $n \geq 0$, $\Omega_n \setminus \Lambda_n$ is a Parseval frame wavelet set, and Ω_n is a frame wavelet set with frame bounds 1 and 2, cf. Proposition 2.2 of [39] (2002).*

Proof. By Proposition 2.1, Lemma 2.1, and the inclusion $\Omega_n \setminus \Lambda_n \subseteq \Omega_n$, it is clear that $\Omega_n \setminus \Lambda_n$ is a Parseval frame wavelet set.

Let $\widehat{\psi} = \mathbf{1}_{\Omega_n}$. Then,

$$\sum_{j \in \mathbb{Z}} |\widehat{\psi}(2^j \xi)|^2 = \sum_{j \in \mathbb{Z}} \mathbf{1}_{\Omega_n \setminus \Lambda_n}(2^j \xi) + \sum_{j \in \mathbb{Z}} \mathbf{1}_{\Lambda_n}(2^j \xi) = 1 + \mathbf{1}_{\bigcup_{j \in \mathbb{Z}} 2^{-j} \Lambda_n}(\xi).$$

It is straightforward from the definition that the sets $2^{-j} \Lambda_n$, $j \in \mathbb{Z}$, are mutually disjoint. This justifies the second equation. Therefore, $\sup_{\xi \in \widehat{\mathbb{R}}^d} \sum_{j \in \mathbb{Z}} |\widehat{\psi}(2^j \xi)|^2 = 2$ and $\inf_{\xi \in \widehat{\mathbb{R}}^d} \sum_{j \in \mathbb{Z}} |\widehat{\psi}(2^j \xi)|^2 = 1$. Since Ω_n tiles $\widehat{\mathbb{R}}^d$ by \mathbb{Z}^d -translation, we have $\widehat{\psi}(2^j \xi) \widehat{\psi}(2^j \xi + k) = \mathbf{1}_{\Omega_n}(2^j \xi) \mathbf{1}_{\Omega_n - k}(2^j \xi) = 0$ for all $j \in \mathbb{Z}$ and $k \in \mathbb{Z}^d \setminus \{0\}$. Hence, we can invoke Theorem 2.8 to assert that Ω_n is a frame wavelet set with frame bounds 1 and 2. \square

2.5 Examples of Euclidean Wavelet Sets

Theorem 2.10. *a. A convex set C cannot partition $\widehat{\mathbb{R}}^d$ by dilations; that is, the set $\{2^j C : j \in \mathbb{Z}\}$ cannot be a tiling of $\widehat{\mathbb{R}}^d$.
b. A convex set is not a wavelet set.*

Proof. a. i. The statement is clear in $\widehat{\mathbb{R}}$, and, for simplicity, we shall only prove the result in the case $d = 2$. The proof can be easily generalized to $\widehat{\mathbb{R}}^d$ by replacing lines in $\widehat{\mathbb{R}}^2$ with hyperplanes in $\widehat{\mathbb{R}}^d$.

ii. Suppose that C is a convex set and that $\{2^j C : j \in \mathbb{Z}\}$ is a partition of $\widehat{\mathbb{R}}^2$. Define

$$S = \left\{ \gamma \in \widehat{\mathbb{R}}^2 : \gamma \in C \text{ and } -\gamma \in C \right\} = C \cap (-C) \subseteq C.$$

It is clear from the definition that the set S is symmetric, i.e., $\gamma \in S$ if and only if $-\gamma \in S$. Clearly, S is also convex. Most important, $|S| = 0$. This is proved by assuming $|S| > 0$ allowing us to verify that C must contain a neighborhood of 0, which, in turn, contradicts the disjointness of $2^j C$, $j \in \mathbb{Z}$, thereby giving $|S| = 0$.

iii. Next, we note that

$$(2^j C) \cap (-2^i C) \subseteq \begin{cases} 2^j S & \text{if } j \geq i \\ 2^i S & \text{if } j \leq i. \end{cases} \quad (2.13)$$

Thus, with our assumption $\bigcup_{j \in \mathbb{Z}} 2^j C = \widehat{\mathbb{R}}^2$, we compute

$$\begin{aligned} |\widehat{\mathbb{R}}^2| &= \left| \left(\bigcup_{j \in \mathbb{Z}} 2^j C \right) \cap \left(\bigcup_{i \in \mathbb{Z}} -2^i C \right) \right| \\ &= \left| \bigcup_{i, j \in \mathbb{Z}} 2^j C \cap (-2^i C) \right| \\ &\leq \sum_{i, j \in \mathbb{Z}} |2^{\max(i, j)} S| = 0 \end{aligned} \quad (2.14)$$

by (2.13) and since $|S| = 0$. (2.14) is obviously false, and so the proof is complete.

b. This is immediate from part a and Theorem 2.1. \square

Theorem 2.11. *Let C_1, C_2, \dots, C_n be convex sets in $\widehat{\mathbb{R}}^d$. If*

$$\dot{\bigcup}_{i=1}^n \bigcup_{j \in \mathbb{Z}} 2^j C_i = \widehat{\mathbb{R}}^d,$$

where the left side is a disjoint union, then $n \geq d + 1$. In particular, if C_1, C_2, \dots, C_n are convex sets in $\widehat{\mathbb{R}}^d$ such that $\bigcup_{i=1}^n C_i$ is a wavelet set, then $n \geq d + 1$, see [23] (2006).

By definition, $W \subseteq \widehat{\mathbb{R}}^d$ is a *subspace wavelet set* if the family

$$\left\{ (\mathbf{1}_W^\vee)_{m,n} : m \in \mathbb{Z}, n \in \mathbb{Z}^d \right\}$$

is an ONB for a subspace of $L^2(\mathbb{R}^d)$.

Theorem 2.12. *The set $W \subseteq \widehat{\mathbb{R}}^d$ is a subspace wavelet set if and only if the characteristic function $\mathbf{1}_E$ of the set $E = \bigcup_{j < 0} 2^j W$ satisfies the following consistency equation:*

$$1 + \sum_{k \in \mathbb{Z}^d} \mathbf{1}_E(\gamma + k) = \sum_{k \in \mathbb{Z}^d} \mathbf{1}_E\left(\frac{1}{2}(\gamma + k)\right) \quad \text{a.e.}$$

In particular, W is a wavelet set for all of $\widehat{\mathbb{R}}^d$ if and only if, in addition, $\bigcup_{j \in \mathbb{Z}} 2^j E$ contains, up to a set of measure zero, a neighborhood of the origin, see [23] (2006).

Theorem 2.13. *Every wavelet set in $[-1, 1]^d \setminus [-\frac{1}{4}, \frac{1}{4}]^d$ can be constructed by means of the NMC.*

The proof has some intricacies.

Corollary 2.2. *For every wavelet set Ω in $[-1, 1]^d$,*

$$\left| \Omega \cap [-\alpha, \alpha]^d \right| > 0 \quad \text{for all } \alpha > \frac{1}{4}.$$

Proof. If $\left| \Omega \cap [-\frac{1}{4}, \frac{1}{4}]^d \right| > 0$, then we are done. Otherwise, by Theorem 2.13, the wavelet set Ω can be constructed by the NMC. By definition of the mapping T ,

it can be shown that Λ_2 is the only one of the sets Λ_n that intersects the “square annulus” $[-\frac{5}{16}, \frac{5}{16}]^d \setminus [-\frac{1}{4}, \frac{1}{4}]^d$. Again, since T is a \mathbb{Z}^d -translated mapping, it is not possible for Λ_2 to cover $[-\alpha, \alpha]^d \setminus [-\frac{1}{4}, \frac{1}{4}]^d$ for any $\alpha > \frac{1}{4}$. This completes the proof. \square

Theorem 2.14. *For any $\alpha < 1$, a wavelet set Ω can not be contained in $[-\alpha, \alpha]^d$.*

Proof. Suppose that $\Omega \subseteq [-\alpha, \alpha]^d$ is a wavelet set. Then the integral translates of Ω will tile $\widehat{\mathbb{R}}^d$, i.e.,

$$\dot{\bigcup}_{k \in \mathbb{Z}^d} (\Omega + k) = \widehat{\mathbb{R}}^d,$$

where $\dot{\bigcup}$ designates disjoint union.

Observe that for each fixed $i = 1, \dots, d$, the union of all translates $\Omega + (k_1, \dots, k_d)$ with $k_i \neq 0$ leaves out the band $B_i = \{(x_1, x_2, \dots, x_d) \in \mathbb{R}^d : \alpha - 1 < x_i < 1 - \alpha\}$, i.e.,

$$\dot{\bigcup}_{k_i \neq 0} (\Omega + k) \cap B_i = \emptyset, \text{ for each } i = 1, 2, \dots, d.$$

Therefore,

$$\dot{\bigcup}_{k \neq (0, \dots, 0)} (\Omega + k) \cap B = \emptyset,$$

where $B = \bigcap_{i=1}^d B_i = (\alpha - 1, 1 - \alpha)^d$. This clearly implies that $B \subseteq \Omega$ and that $\emptyset \neq B \subseteq \Omega \cap 2\Omega$, a contradiction to the dyadic-dilation congruence property of wavelet sets. \square

Example 2.1. Figure 2.1 is the $\Omega_7 \subseteq \widehat{\mathbb{R}}^2$ approximant of the 2-dimensional generalization of the Shannon wavelet set described in Sect. 2.2.2 for the “variables” $\Omega_0 = Q$ and $T(\gamma_1, \gamma_2) = (\gamma_1, \gamma_2) - (\text{sign}(\gamma_1), \text{sign}(\gamma_2))$.




Fig. 2.1 The set Ω_7 for the 2-dimensional Shannon wavelet set of Example 2.1

Example 2.2. Wedding sets

The wedding cake wavelet set was defined in [34] (1998), [36] (1998), and it can be constructed by our NMC method with $\Omega_0 = \mathbb{Q}$ and $T(\gamma_1, \gamma_2) = (\gamma_1 + \text{sign } \gamma_1, \gamma_2)$ for $(\gamma_1, \gamma_2) \in \Omega_0$. It was introduced as a simple wavelet set consisting of 3 connected sets, see Fig. 2.2. The interior of each component is connected. We constructed the wedding night set in Fig. 2.3, and it consists of two connected sets. The NMC also allowed us to construct an alternative wedding cake set consisting of two connected sets with connected interiors (Fig. 2.4).

Example 2.3. An example of a connected wavelet set, in which the interior consists of infinitely many components, was given in [21], see Fig. 2.5. It is believed that there is no connected wavelet set with connected interior.




Fig. 2.2 The wedding cake set of Example 2.2




Fig. 2.3 The wedding night set of Example 2.2

Fig. 2.4 The wedding cake set of Example 2.2 consisting of 2 connected sets





Fig. 2.5 The connected wavelet set of Example 2.3



Example 2.4. A generalization of the Journé wavelet set.


Since the Journé wavelet set can be constructed by the NMC with $\Omega_0 = [-\frac{1}{2}, \frac{1}{2})$ and $T(\gamma) = \gamma + 2 \operatorname{sign}(\gamma)$, one of its d -dimensional versions can be produced by setting $\Omega_0 = Q$ and

$$T(\gamma_1, \dots, \gamma_d) = (\gamma_1 + 2 \operatorname{sign}(\gamma_1), \dots, \gamma_d + 2 \operatorname{sign}(\gamma_d)),$$

see Fig. 2.6 for the corresponding wavelet set in $\widehat{\mathbb{R}}^2$.

It should be noted that the NMC only produces wavelet sets that are bounded away from the origin and infinity, i.e., they have holes at the origin and are bounded sets. Related work can be found in [73] (2002).

Fig. 2.6 The 2-dimensional Journé wavelet set of Example 2.4



2.6 Locally Compact Abelian Groups, Wavelets, and the p -adic Field

2.6.1 The p -adic Field \mathbb{Q}_p

Theorem 2.15. *Every locally compact abelian group (LCAG) G is topologically and algebraically isomorphic to $\mathbb{R}^d \times G_o$, where \mathbb{R}^d is Euclidean space and G_o is a LCAG containing a compact open subgroup H_o .*

This fact follows from a result in [1] (1965) combined with [56, Sect. 9.8] (1963). For the remainder, we shall deal with wavelet theory for functions defined on groups G_o .

A common example of a group G_o is the field \mathbb{Q}_p of p -adic rationals. Given any prime number p , the field \mathbb{Q}_p is the completion of the field \mathbb{Q} of rationals with respect to the p -adic absolute value $|p^r m/n|_p = p^{-r}$ for all $r, m, n \in \mathbb{Z}$ with m and n not divisible by p . Equivalently, \mathbb{Q}_p may be thought of as the set of Laurent series in the “variable” p , with coefficients $0, 1, \dots, p$. This means that

$$\mathbb{Q}_p = \left\{ \sum_{n \geq n_0} a_n p^n : n_0 \in \mathbb{Z} \text{ and } a_n \in \{0, 1, \dots, p-1\} \right\},$$

with addition and multiplication as usual for Laurent series, except with carrying of digits, so that, e.g., in \mathbb{Q}_7 , we have

$$(4 + 1 \cdot 7) + (6 + 5 \cdot 7) = 3 + 0 \cdot 7 + 1 \cdot 7^2.$$

The p -adic absolute value extends naturally to \mathbb{Q}_p , and under the operation of addition, \mathbb{Q}_p forms a LCAG, with topology induced by $|\cdot|_p$, and with compact open subgroup \mathbb{Z}_p , the ring of p -adic integers, consisting of Taylor series in p . Equivalently, \mathbb{Z}_p is the closure of $\mathbb{Z} \subseteq \mathbb{Q}_p$ with respect to $|\cdot|_p$. Further, $\mathbb{Z}_p = \{x \in \mathbb{Q}_p : |x|_p \leq 1\}$.




Fig. 2.7 \widehat{G} for a LCAG G with compact open subgroup H and expansive automorphism A , with $|A| = 3$

A few other details about \mathbb{Q}_p appear in Sect. 2.6.3. In addition, Sect. 2.7.2 contains an informal discussion of the geometry of \mathbb{Q}_p and related groups, including a rough sketch of \mathbb{Q}_3 in Fig. 2.7.

2.6.2 Wavelet Theories on \mathbb{Q}_p and Related Groups

The main obstacle to producing a theory of wavelets on groups like \mathbb{Q}_p is that \mathbb{Q}_p has no nontrivial discrete subgroups, and thus there is no lattice to use for translations. Some other LCAGs have both a compact open subgroup and a discrete cocompact lattice. For example, Lang [69] (1996), [70] (1998), [71] (1998) constructed wavelets for the Cantor dyadic group (known to number theorists as the field $\mathbb{F}_2((t))$ of formal Laurent series over the field of two elements) using the lattice consisting of polynomials in t^{-1} with trivial constant term. Farkov [45] (2005), [44] (2005) later generalized Lang's construction to other LCAGs with a compact open subgroup and a discrete cocompact lattice. However, for \mathbb{Q}_p and other LCAGs with compact open subgroups, the lack of a lattice requires a different strategy.

Several authors [65] (2002), [59] (2004), [62] (2009), [90] (2009) have constructed wavelets on such groups by the following strategy. Given a LCAG G with compact open subgroup H , choose a set \mathcal{C} of coset representatives for G/H , and translate only by elements of \mathcal{C} . For example, if $G = \mathbb{Q}_p$ and $H = \mathbb{Z}_p$, we may choose \mathcal{C} to consist of all elements of \mathbb{Q}_p of the form a/p^n , where $n \geq 1$ and $0 \leq a \leq p^n - 1$, so that every element of \mathbb{Q}_p may be represented uniquely as $x + s$, for $x \in \mathbb{Z}_p$ and $s \in \mathcal{C}$. Then, given an appropriate dilation operator $A : G \rightarrow G$ (such as multiplication-by- $1/p$, in the case of \mathbb{Q}_p), it is possible to develop a corresponding wavelet theory. For example, the p -adic wavelets of [65], as well as of those on

certain non-group ultrametric spaces in [61] (2005), are simply inverse transforms of characteristic functions of disks, and they were shown to be wavelets by direct computation. Meanwhile, the wavelets of [59], [62], and [90] all arise from a p -adic version of MRA. Common to all of the above constructions, however, is the use of translations by coset representatives.

Unfortunately, the chosen set of coset representatives is usually not a group; e.g., the set $\mathcal{C} \subseteq \mathbb{Q}_p$ of the previous paragraph is closed neither under addition nor under additive inverses. As a consequence, the resulting theory seems limited. In particular, all such wavelets known to date are step functions with a finite number of steps.

Instead, we present a different wavelet theory for such groups, using a different set of operators in place of translation by elements of \mathcal{C} . Although our operators are not actual translations, they have the crucial advantage of forming a group. We call them *pseudo-translations*. The resulting theory allows a much wider variety of wavelets, including most of the wavelets produced by other authors, as well as many others. After presenting the general theory and definition of our wavelets, we shall show that it is possible to construct many such wavelets using a theory of wavelet sets, and we shall give an algorithm for constructing a wide variety of wavelet sets. We expect that it should also be possible to develop multiresolution analysis for our wavelet theory, but this has not yet been done.

2.6.3 Prerequisites About LCAGs

In this section, we set some notation and recall a few standard facts about abstract LCAGs; see [56] (1963), [57] (1970), [86] (1966), [87] (1968), and [88] (1962) for details.

Let G be a LCAG with compact open subgroup H . Denote by \widehat{G} the dual group of G , with action denoted $(x, \gamma) \in \mathbb{C}^\times$, for $x \in G$ and $\gamma \in \widehat{G}$. The *annihilator subgroup* of H in \widehat{G} is

$$H^\perp = \{\gamma \in \widehat{G} : \forall x \in H, (x, \gamma) = 1\} \subseteq \widehat{G},$$

which is, in turn, a compact open subgroup of \widehat{G} .

The quotient group G/H of course consists of cosets $x + H$, also denoted $[x]$, for $x \in G$. This quotient is discrete, because H is open in G . Moreover, G/H is isomorphic as a LCAG to the dual of H^\perp . The isomorphism is easy to write down; the element $x + H \in G/H$ acts on H^\perp by $(x + H, \gamma) = (x, \gamma)$, for any $\gamma \in H^\perp$. Similarly, \widehat{H} and \widehat{G}/H^\perp are isomorphic discrete groups.

Set $\mu = \mu_G$ and $\nu = \nu_{\widehat{G}}$ to be Haar measures on G and \widehat{G} , respectively, normalized so that $\mu(H) = \nu(H^\perp) = 1$. These normalizations induce counting measures on the discrete groups G/H and \widehat{G}/H^\perp , and they make the Fourier transform, given by

$$\widehat{f}(\gamma) = \int_G f(x) \overline{(x, \gamma)} d\mu(x), \quad \text{for all } f \in L^2(G),$$

an isometry between $L^2(G)$ and $L^2(\widehat{G})$. See, for instance, [57, Sect. 31.1] (1970), [87] (1968), and [17, Sect. 1.3] (2004).

By way of example, consider again the case $G = \mathbb{Q}_p$ and $H = \mathbb{Z}_p$. The quotient $\mathbb{Q}_p/\mathbb{Z}_p$ is isomorphic to μ_{p^∞} , the subgroup of \mathbb{C}^\times consisting of all roots of unity ζ for which $\zeta^{p^n} = 1$ for some $n \geq 0$. Meanwhile, \mathbb{Q}_p is self-dual, with duality action given by $(x, \gamma) = \chi(x\gamma)$, where $\chi : \mathbb{Q}_p \rightarrow \mathbb{C}$ is the character given by

$$\chi \left(\sum_{n \geq n_0} a_n p^n \right) = \exp \left(2\pi i \sum_{n=n_0}^{-1} a_n p^n \right).$$

The annihilator \mathbb{Z}_p^\perp is just \mathbb{Z}_p under this self-duality.

Our wavelet theory will of course require a dilation operator. Given an automorphism $A : G \rightarrow G$, there is a unique positive number $|A|$, the *modulus* of A , with the property that for any measurable set $U \subseteq G$, we have $\mu(AU) = |A|\mu(U)$. Therefore, for any $f \in C_c(G)$, $\int_G f \circ A(x) d\mu(x) = |A|^{-1} \int_G f(x) d\mu(x)$. See, e.g., [56, Sect. 15.26] (1963). In addition, A has an adjoint element $A^* : \widehat{G} \rightarrow \widehat{G}$, defined by $(Ax, \gamma) = (x, A^* \gamma)$ for all $x \in G$ and $\gamma \in \widehat{G}$. We have $(A^*)^{-1} = (A^{-1})^*$, $|A|^{-1} = |A^{-1}|$, and $|A^*| = |A|$.

2.7 Wavelets for Groups with Compact Open Subgroups

2.7.1 Pseudo-Translations

In this section, we present the pseudo-translation operators to be used in our wavelet theory. Rather than translating by one fixed element of each coset $[s] \in G/H$, we shall construct an operator $\tau_{[s]} : L^2(G) \rightarrow L^2(G)$ for each $[s] \in G/H$ determined only by the coset $[s] = s + H$, and not by a choice of a particular coset representative $s_0 \in [s]$. In addition, our operators will form a group, in that $\tau_{[s]+[t]} = \tau_{[s]}\tau_{[t]}$. The resulting operators are usually not true translations, but $\tau_{[s]}$ will still be similar in certain ways to the translation-by- s operator.

To construct our operators, however, we *shall* have to make a choice of coset representative; but we choose a set \mathcal{D} of coset representatives in \widehat{G} for \widehat{G}/H^\perp , rather than representatives in G for G/H . That is, $\mathcal{D} \subseteq \widehat{G}$ is a discrete subset (probably not forming a subgroup) consisting of exactly one element of every coset $\sigma + H^\perp$. We then define $\tau_{[s]}$ by its induced dual map $\widehat{\tau}_{[s]} : L^2(\widehat{G}) \rightarrow L^2(\widehat{G})$, as follows.

Definition 2.4. Let G be a LCAG with compact open subgroup $H \subseteq G$. Let $\mathcal{D} \subseteq \widehat{G}$ be a set of coset representatives in \widehat{G} for the quotient $\widehat{H} = \widehat{G}/H^\perp$.

Define the map $\eta = \eta_{\mathcal{D}} : \widehat{G} \rightarrow H^\perp \subseteq \widehat{G}$ by

$$\eta(\gamma) = \text{the unique } \beta \in H^\perp \text{ such that } \gamma - \beta \in \mathcal{D}.$$

For each $[s] \in G/H$, the pseudo-translation-by- $[s]$ operator $\tau_{[s]} = \tau_{[s], \mathcal{D}} : L^2(G) \rightarrow L^2(G)$ is given by

$$\widehat{\tau_{[s]}f}(\gamma) = \overline{(s, \eta_{\mathcal{D}}(\gamma))} \widehat{f}(\gamma).$$

Note that the true translation-by- s operator $T_s : L^2(G) \rightarrow L^2(G)$ acts on the transform side by

$$\widehat{T_s f}(\gamma) = \overline{(s, \gamma)} \widehat{f}(\gamma).$$

Thus, $\tau_{[s]}$ resembles a translation operator except for the correction by $\eta_{\mathcal{D}}$. The function $\eta_{\mathcal{D}}(\gamma)$, in turn, should be viewed as giving the difference between γ and the nearest “lattice” point, where we consider \mathcal{D} to be an analog of the dual lattice. In the Euclidean setting, where \mathcal{D} really is a dual lattice and the translating element s really is in a lattice, the corresponding quantity $(s, \eta_{\mathcal{D}}(\gamma))$ would exactly equal (s, γ) . Thus, $\eta_{\mathcal{D}}$ should be thought of as correcting for the fact that \mathcal{D} is not actually a lattice.

The following proposition shows that the other promised properties of $\tau_{[s]}$ also hold.

Proposition 2.2. *Let G, H , and \mathcal{D} be as in Definition 2.4. We have the following.*

- a. $\tau_{[s], \mathcal{D}}$ is well defined, i.e., if $s + H = t + H$, then $\tau_{[s], \mathcal{D}} = \tau_{[t], \mathcal{D}}$.
- b. $\tau_{[0], \mathcal{D}} f = f$ for all $f \in L^2(G)$.
- c. $\tau_{[s], \mathcal{D}} \circ \tau_{[t], \mathcal{D}} = \tau_{[s+t], \mathcal{D}}$ for all $s, t \in G$.

Proof. Given any $s, t \in G$ lying in the same coset $s + H = t + H$ and any $\gamma \in \widehat{G}$, we have

$$(t, \eta_{\mathcal{D}}(\gamma)) = (s, \eta_{\mathcal{D}}(\gamma)) (t - s, \eta_{\mathcal{D}}(\gamma)) = (s, \eta_{\mathcal{D}}(\gamma)),$$

because $t - s \in H$ and $\eta_{\mathcal{D}}(\gamma) \in H^\perp$. Part *a* follows. Similarly, parts *b* and *c* are immediate from the observations that $(0, \eta_{\mathcal{D}}(\gamma)) = 1$ and $(s + t, \eta_{\mathcal{D}}(\gamma)) = (s, \eta_{\mathcal{D}}(\gamma)) (t, \eta_{\mathcal{D}}(\gamma))$. \square

Besides the elegant properties listed in Proposition 2.2, the reason for the particular forms of $\tau_{[s], \mathcal{D}}$ and $\eta_{\mathcal{D}}$ will become clear in (2.20), during the proof of Theorem 2.16.

2.7.2 Expansive Automorphisms and Dilations

When constructing wavelets in $L^2(\mathbb{R}^d)$, one cannot use just any automorphism $A : G \rightarrow G$ as a dilation operator, but rather one with particular properties with respect to the lattice. We now present the corresponding property needed for dilations in our setting.

Definition 2.5. Let G be a LCAG with compact open subgroup $H \subseteq G$, and let $A : G \rightarrow G$ be an automorphism. We say that A is *expansive* with respect to H if both of the following conditions hold:

- i. $H \subsetneq AH$, and
- ii. $\bigcap_{n \leq 0} A^n H = \{0\}$.

As noted in [17, Sect. 2.2] (2004), if G has a compact open subgroup H and expansive automorphism A , then $|A|$ is an integer strictly greater than 1, G/H is infinite, and G is not compact. In addition, on the dual side, we have $H^\perp \subsetneq A^*H^\perp$, and $\bigcup_{n \geq 0} A^{*n}H^\perp = \widehat{G}$.

The expansiveness condition, together with the original assumption that G has a compact open subgroup, says that G and \widehat{G} both have a self-similar structure. In particular, if we sketch H^\perp as a disk, then \widehat{G} is a union of larger and large dilates of that disk. Meanwhile, each dilate $A^{*n}H^\perp$ contains finitely many (in fact, exactly $|A|^n$) translates (i.e., cosets $\sigma + H^\perp$) of H^\perp . Similarly, applying negative powers of A^* , we can see that H^\perp itself consists of $|A|$ translates of the smaller disk $(A^*)^{-1}H^\perp$, each of which itself consists of $|A|$ translates of the still smaller disk $(A^*)^{-2}H^\perp$, and so on. Thus, H^\perp has a fractal structure, much like the Cantor set, while \widehat{G} is an infinite union of translates of H^\perp . See Fig. 2.7 for a sketch of such a group \widehat{G} with an expansive automorphism of modulus 3.

For example, if $G = \mathbb{Q}_p$ and $H = \mathbb{Z}_p$, we may choose $A : \mathbb{Q}_p \rightarrow \mathbb{Q}_p$ to be $A(x) = x/p$, which maps \mathbb{Z}_p to $(1/p)\mathbb{Z}_p \subsetneq \mathbb{Z}_p$, satisfying condition *i* of Definition 2.5. Condition *ii* also holds, because $\bigcap_{n \leq 0} p^n \mathbb{Z}_p = \{0\}$. The modulus in this case is $|A| = |1/p|_p = p$. Figure 2.7 may therefore be considered to be a rough sketch of \mathbb{Q}_3 .

2.7.3 Wavelets

As in the Euclidean setting, an automorphism $A : G \rightarrow G$ induces an operator on $L^2(G)$, sending $f(x)$ to $|A|^{1/2}f(Ax)$; the constant in front, of course, ensures that the resulting operator is unitary. Thus, we may make the following definition.

Definition 2.6. Let G be a LCAG with compact open subgroup $H \subseteq G$, let \mathcal{D} be a choice of coset representatives in \widehat{G} for $\widehat{H} = \widehat{G}/H^\perp$, let $A : G \rightarrow G$ be an automorphism, and consider $[s] \in G/H$. The *dilated translate* of $f \in L^2(G)$ is defined to be

$$f_{A,[s]}(x) = |A|^{1/2} \cdot (\tau_{[s],\mathcal{D}}f)(Ax). \quad (2.15)$$

Note that (2.15) implies that

$$\widehat{f_{A,[s]}}(\gamma) = |A|^{-1/2} \widehat{f}((A^*)^{-1}\gamma) \overline{(s, \eta((A^*)^{-1}\gamma))}. \quad (2.16)$$

Now that we have appropriate dilation and translation operators, we are prepared to define wavelets on our group G .

Definition 2.7. Let G be a LCAG with compact open subgroup $H \subseteq G$, let $\mathcal{D} \subseteq \widehat{G}$ be a choice of coset representatives in \widehat{G} for \widehat{G}/H^\perp , and let $A : G \rightarrow G$ be an automorphism. Consider $\Psi = \{\psi_1, \dots, \psi_N\} \subseteq L^2(G)$. We say Ψ is a *set of wavelet generators* for $L^2(G)$ with respect to \mathcal{D} and A if

$$\{\psi_{j,m,[s]} : 1 \leq j \leq N, m \in \mathbb{Z}, [s] \in G/H\}$$

forms an ONB for $L^2(G)$, where

$$\psi_{j,m,[s]}(x) = |A|^{m/2} \cdot (\tau_{[s],\mathcal{D}}\psi_j)(A^m x),$$

as in (2.15). In that case, the resulting basis is called a *wavelet basis* for $L^2(G)$.

If $\Psi = \{\psi\}$, then ψ is a *single wavelet* for $L^2(G)$.

2.8 Geometry of Wavelet Sets for G

As we did for $L^2(\mathbb{R}^d)$, we shall use the machinery of wavelet sets, and not MRA, to construct wavelets for $L^2(G)$. Therefore, we state the following definition, cf. [34] (1998), [35] (1997).

Definition 2.8. Let G , H , \mathcal{D} , and A be as in Definition 2.7. Let $\Omega_1, \dots, \Omega_N$ be measurable subsets of \widehat{G} , and let $\psi_j = \mathbf{1}_{\Omega_j}^\vee$ for each $j = 1, \dots, N$. We say that $\{\Omega_1, \dots, \Omega_N\}$ is a *wavelet collection of sets* if $\Psi = \{\psi_1, \dots, \psi_N\}$ is a set of wavelet generators for $L^2(G)$.

If $N = 1$, then $\Omega = \Omega_1$ is a *wavelet set*.

We shall characterize wavelet sets in terms of properties analogous to the Euclidean notions of τ -congruence and δ -congruence, as described in Sect. 2.2. See also [17, Sect. 3.2] (2004) for a broader discussion in our setting.

Definition 2.9. Let G be a LCAG with compact open subgroup $H \subseteq G$, let $\mathcal{D} \subseteq \widehat{G}$ be a choice of coset representatives in \widehat{G} for $\widehat{H} = \widehat{G}/H^\perp$, and let $\Omega \subseteq \widehat{G}$ be a subset. We say Ω is (τ, \mathcal{D}) -congruent to H^\perp if there exist measure zero subsets $V_0 \subseteq \Omega$ and $V'_0 \subseteq H^\perp$, a sequence $\{\sigma_n\}_{n \geq 1} \subseteq \mathcal{D}$, and a countable partition $\{V_n : n \geq 1\}$ of $\Omega \setminus \{V_0\}$ into measurable subsets such that $\{V_n - \sigma_n : n \geq 1\}$ forms a partition of $H^\perp \setminus V'_0$.

Definition 2.10. Let $\{W_m : m \in \mathbb{Z}\}$ be a countable set of measurable subsets of \widehat{G} . We say that $\{W_m\}$ *tiles* \widehat{G} if

$$v\left(\widehat{G} \setminus \left[\bigcup_{m \in \mathbb{Z}} W_m\right]\right) = 0$$

and

$$v(W_m \cap W_n) = 0, \quad \text{for all } m, n \in \mathbb{Z}, m \neq n.$$

Our first main result characterizes wavelet collections of sets in terms of the two preceding definitions.

Theorem 2.16. Let G be a LCAG with compact open subgroup $H \subseteq G$, let $\mathcal{D} \subseteq \widehat{G}$ be a choice of coset representatives in \widehat{G} for \widehat{G}/H^\perp , and let $A : G \rightarrow G$

be an automorphism. A finite set $\{\Omega_1, \dots, \Omega_N\}$ of measurable subsets of \widehat{G} is a wavelet collection of sets if and only if both of the following conditions hold:

- i. $\{A^{*n}\Omega_j : n \in \mathbb{Z}, j = 1, \dots, N\}$ tiles \widehat{G} , and
- ii. $\forall j = 1, \dots, N$, Ω_j is (τ, \mathcal{D}) -congruent to H^\perp .

In that case, \widehat{G} is σ -compact, each $v(\Omega_j) = 1$, and each $\mathbf{1}_{\Omega_j} \in L^2(\widehat{G})$.

Proof. See [17, Theorem 3.4]. The centerpiece of the proof is to show that

$$\sum_{j=1}^N \sum_{m \in \mathbb{Z}} \sum_{[s] \in G/H} |\langle f, \psi_{j,m,[s]} \rangle|^2 = \|f\|_2^2, \quad \text{for all } f \in L^2(G), \quad (2.17)$$

at least under the assumptions that the sum on left side of (2.17) converges (and in particular, all but countably many terms of the sum are 0), and that properties i and ii of Theorem 2.16 hold, cf. the calculation (2.7) in Sect. 2.2.2. We now reproduce the argument from [17].

By Plancherel's theorem and (2.16), we have

$$\begin{aligned} \sum_{j,m,[s]} |\langle f, \psi_{j,m,[s]} \rangle|^2 &= \sum_{j,m,[s]} |\langle \widehat{f}, \widehat{\psi_{j,m,[s]}} \rangle|^2 \\ &= \sum_{j,m,[s]} |A|^{-m} \left| \int_{\widehat{G}} \widehat{f}(\gamma) \cdot \overline{\widehat{\psi_j((A^*)^{-m}\gamma)}} \cdot (s, \eta((A^*)^{-m}\gamma)) \, dv(\gamma) \right|^2 \\ &= \sum_{j,m,[s]} |A|^m \left| \int_{\Omega_j} \widehat{f}(A^{*m}\beta) \cdot (s, \eta(\beta)) \, dv(\beta) \right|^2, \end{aligned} \quad (2.18)$$

where we have substituted $\beta = (A^*)^{-m}\gamma$. By property ii, each Ω_j is (τ, \mathcal{D}) -congruent to H^\perp , thereby giving us partitions $\{V_{j,n}\}_{n \geq 0}$ of Ω_j with $v(V_{j,0}) = 0$ and sequences $\{\sigma_{j,n}\}_{n \geq 1} \subseteq \mathcal{D}$, as in Definition 2.9. Thus, the right side of (2.18) becomes

$$\begin{aligned} &\sum_{j,m,[s]} |A|^m \left| \sum_{n \geq 1} \left[\int_{V_{j,n}} \widehat{f}(A^{*m}\beta) \cdot (s, \eta(\beta)) \, dv(\beta) \right] \right|^2 \\ &= \sum_{j,m,[s]} |A|^m \left| \sum_{n \geq 1} \left[\int_{V_{j,n} - \sigma_{j,n}} \widehat{f}(A^{*m}(\alpha + \sigma_{j,n})) \cdot (s, \eta(\alpha + \sigma_{j,n})) \, dv(\alpha) \right] \right|^2 \\ &= \sum_{j,m,[s]} |A|^m \left| \sum_{n \geq 1} \left[\int_{\widehat{G}} \mathbf{1}_{V_{j,n} - \sigma_{j,n}}(\alpha) \cdot \widehat{f}(A^*(\alpha + \sigma_{j,n})) \cdot (s, \eta(\alpha + \sigma_{j,n})) \, dv(\alpha) \right] \right|^2, \end{aligned} \quad (2.19)$$

where we have substituted $\alpha = \beta - \sigma_{j,n}$. Since $\alpha \in V_{j,n} - \sigma_{j,n} \subseteq H^\perp$, the unique point in $(\alpha + \sigma_{j,n} + H^\perp) \cap \mathcal{D}$ is $\sigma_{j,n}$, and therefore

$$\eta(\alpha + \sigma_{j,n}) = (\alpha + \sigma_{j,n}) - \sigma_{j,n} = \alpha. \quad (2.20)$$

As noted in Sect. 2.7.1, the convenient simplification of Equation (2.20) helps illustrate the reason for the otherwise peculiar-looking description of η and $\tau_{[s],\mathcal{D}}$ in Definition 2.4.

Next, we claim we can exchange the inner summation and integral signs in the last term of (2.19). After all, we know that $\{V_{j,n} - \sigma_{j,n} : n \geq 1\}$ tiles H^\perp . Hence, denoting the integrand of (2.19) by $F_{j,n}$, writing $F_j = \sum_{n \geq 1} F_{j,n}$, and noting that $F_{j,n}$ vanishes off of $V_{j,n} - \sigma_{j,n}$, we see that $F_{j,n}, F_j \in L^2(H^\perp) \subseteq L^1(H^\perp)$, and therefore

$$\sum_{n \geq 1} \int_{\hat{G}} F_{j,n}(\alpha) \, d\nu(\alpha) = \sum_{n \geq 1} \int_{V_{j,n} - \sigma_{j,n}} F_j(\alpha) \, d\nu(\alpha) = \int_{H^\perp} F_j(\alpha) \, d\nu(\alpha).$$

Thus, the right side of (2.19) becomes

$$\sum_{j,m} |A|^m \sum_{[s] \in G/H} \left| \int_{H^\perp} \left[\sum_{n \geq 1} \mathbf{1}_{V_{j,n} - \sigma_{j,n}}(\alpha) \cdot \hat{f}(A^{*m}(\alpha + \sigma_{j,n})) \right] (s, \alpha) \, d\nu(\alpha) \right|^2. \quad (2.21)$$

Because G/H is the (discrete) dual of H^\perp , Plancherel's theorem tells us that

$$\begin{aligned} \sum_{[s] \in G/H} \left| \int_{H^\perp} g(\alpha)(s, \alpha) \, d\nu(\alpha) \right|^2 &= \sum_{[s] \in G/H} \left| \int_{H^\perp} g(\alpha)([s], \alpha) \, d\nu(\alpha) \right|^2 \\ &= \int_{H^\perp} |g(\alpha)|^2 \, d\nu(\alpha), \end{aligned}$$

for any $g \in L^2(H^\perp)$. Thus, (2.21) becomes

$$\sum_{j,m} |A|^m \int_{H^\perp} \left| \sum_{n \geq 1} \mathbf{1}_{V_{j,n} - \sigma_{j,n}}(\alpha) \hat{f}(A^*(\alpha + \sigma_{j,n})) \right|^2 \, d\nu(\alpha),$$

which, in turn, is

$$\sum_{j,m} |A|^m \int_{H^\perp} \left[\sum_{n \geq 1} |\hat{f}(A^*(\alpha + \sigma_{j,n}))|^2 \mathbf{1}_{V_{j,n} - \sigma_{j,n}}(\alpha) \right] \, d\nu(\alpha), \quad (2.22)$$

because, for fixed f , the sets $V_{j,n} - \sigma_{j,n}$ are pairwise disjoint. We can now interchange the inner summation and integral as before, and (2.22) becomes

$$\begin{aligned} &\sum_{j,m} |A|^m \sum_{n \geq 1} \left[\int_{V_{j,n} - \sigma_{j,n}} |\hat{f}(A^*(\alpha + \sigma_{j,n}))|^2 \, d\nu(\alpha) \right] \\ &= \sum_{j,m} |A|^m \sum_{n \geq 1} \left[\int_{V_{j,n}} |\hat{f}(A^*(\beta))|^2 \, d\nu(\beta) \right] \\ &= \sum_{j,m} |A|^m \int_{\Omega_j} |\hat{f}(A^*(\beta))|^2 \, d\nu(\beta) \\ &= \sum_{j,m} \int_{A^* \Omega_j} |\hat{f}(\gamma)|^2 \, d\nu(\gamma). \end{aligned} \quad (2.23)$$

However, $\{A^*\Omega_j\}$ tiles \widehat{G} . Hence, the right side of (2.23) is

$$\int_{\widehat{G}} |\hat{f}(\gamma)|^2 d\nu(\gamma) = \|\hat{f}\|_2^2 = \|f\|_2^2,$$

proving (2.17). □

2.9 The Construction of Wavelet Sets for G

2.9.1 The Basic Construction

Motivated by the NMC described in Sect. 2.3, we now present an algorithm for constructing wavelet collections of sets. As before, G is a LCAG with compact open subgroup H , \mathcal{D} is a choice of coset representatives in \widehat{G} for the quotient \widehat{G}/H^\perp , and $A : G \rightarrow G$ is an automorphism, which we now assume to be expansive with respect to H .

Our algorithm begins with the following data.

- i. A nonnegative integer $M \geq 0$. Set $W = (A^*)^M H^\perp$.
- ii. A positive integer $N \geq 1$.
- iii. For each $j = 1, \dots, N$, a measurable set $\Omega_{j,0} \subseteq W$ that is (τ, \mathcal{D}) -congruent to H^\perp .
- iv. For each $j = 1, \dots, N$, a measurable injective function $T_j : W \rightarrow (A^*W) \setminus W$ such that

$$T_j(\gamma) = \gamma - \sigma'_j(\gamma) + \sigma_j(\gamma), \quad \text{for all } \gamma \in W,$$

where $\sigma_j(\gamma) \in \mathcal{D}$, and $\sigma'_j(\gamma)$ is the unique element of $\mathcal{D} \cap (\gamma + H^\perp)$.

We also set the following compatibility requirements on the above data.

- v. The union $\widetilde{\Omega}_0 = \bigcup_{j=1}^N \Omega_{j,0}$ contains the neighborhood $(A^*)^{-\ell} H^\perp$ of the origin, for some integer $\ell \geq 0$.
- vi. For any distinct $j, k \in \{1, \dots, N\}$, either

$$T_j W \cap T_k W = \emptyset$$

or

$$T_j = T_k \quad \text{and} \quad \Omega_{j,0} \cap \Omega_{k,0} = \emptyset.$$

Note, however, that we do *not* require the sets $\Omega_{1,0}, \dots, \Omega_{N,0}$ to be disjoint. The possibility that two or more of them overlap will be dealt with in the algorithm to be described below. Also, note that, because A is expansive, the set W contains H^\perp , and $(A^*)^{-1}W$ properly contains W .

Meanwhile, as in Sect. 2.3, the mappings T_j should be understood as slicing W into finitely many measurable pieces and then translating each piece, with the injectivity condition requiring that the images of the pieces do not overlap. In Sect. 2.3,

the translation is by an element of the lattice. In our setting, however, the translation is by an element of the form $\sigma - \sigma'$, where $\sigma, \sigma' \in \mathcal{D}$, and $\sigma' + H^\perp$ contains the piece in question, while $\sigma + H^\perp$ contains its image. This more complicated description is required for the proof of the algorithm's validity; see Sect. 2.9.2.

Algorithm 2.3. Given the initial data described above, our algorithm proceeds inductively, building sets $\Lambda_{j,n+1}$ and $\Omega_{j,n}$ for each $n \geq 0$, as follows. Given the sets $\Omega_{j,n}$ and their union $\tilde{\Omega}_n = \bigcup_{j=1}^n \Omega_{j,n}$ for a particular $n \geq 0$, define $\Lambda_{j,n+1}$ to be the overlap

$$\Lambda_{j,n+1} = \Omega_{j,n} \cap \left[\bigcup_{m \geq 1} (A^*)^{-m} \tilde{\Omega}_n \right]$$

if $n \geq 1$, or

$$\Lambda_{j,1} = \Omega_{j,0} \cap \left[\left(\bigcup_{m \geq 1} (A^*)^{-m} \tilde{\Omega}_n \right) \cup \left(\bigcup_{k=1}^{j-1} \Omega_{k,0} \right) \right]$$

if $n = 0$.

This additional complication at the $n = 0$ step could just as well have been used in the Euclidean setting of \mathbb{R}^d , but it first appeared in [17] (2004) in the non-Euclidean setting in order to give the resulting algorithm the flexibility required to generate certain wavelets previously constructed by Kozyrev [65] (2002).

Then, for each j , build $\Omega_{j,n+1}$ from $\Omega_{j,n}$ by translating $\Lambda_{j,n+1} \subseteq \Omega_{j,n+1}$ by T_j , i.e.,

$$\Omega_{j,n+1} = (\Omega_{j,n} \setminus \Lambda_{j,n+1}) \cup T_j \Lambda_{j,n+1}.$$

Finally, for each $j = 1, \dots, N$, we set

$$\Lambda_j = \bigcup_{m \geq 1} \Lambda_{j,m} \quad \text{and} \quad \Omega_j = (\Omega_{j,0} \setminus \Lambda_j) \cup T_j \Lambda_j. \quad (2.24)$$

Intuitively, Ω_j is a sort of limit of the sequence of sets $\{\Omega_{j,n}\}_{n \geq 0}$. We refer the reader to [17, Sect. 4.1] (2004) for a more detailed description of the algorithm, including verification that $\Lambda_{j,n+1}$ does indeed always lie in W , and hence it makes sense to consider $T_j \Lambda_{j,n+1}$.

2.9.2 Validity of the Algorithm

The following theorem appeared as [17, Theorem 4.2] (2004).

Theorem 2.17. *Let G be a LCAG with compact open subgroup $H \subseteq G$, let $\mathcal{D} \subseteq \hat{G}$ be a choice of coset representatives in \hat{G} for \hat{G}/H^\perp , and let $A : G \rightarrow G$ be an expansive automorphism. Given the data listed in Sect. 2.9.1, the sets $\{\Omega_1, \dots, \Omega_N\}$ of (2.24) produced by the algorithm of Sect. 2.9.1 form a wavelet collection of sets.*

We refer the reader to [17, Sect. 4.2] (2004) for the proof. The idea of the proof is to verify that $\{\Omega_1, \dots, \Omega_N\}$ satisfy conditions *i* and *ii* of Theorem 2.16.

To verify condition *i*, we first check that $\bigcup_{m \in \mathbb{Z}} A^{*m} \tilde{\Omega}$ covers \hat{G} , where $\tilde{\Omega} = \bigcup_{j=1}^N \Omega_j$. This fact follows from the expansiveness of A and the stipulation in the algorithm's initial data that $\tilde{\Omega}_0 \supseteq (A^*)^{-\ell} H^\perp$. To prove that the covering of \hat{G} is in fact a tiling, we first note that $\Omega_{1,1}, \dots, \Omega_{N,1}$ are pairwise disjoint, essentially by definition, because $\Lambda_{j,1}$ contains any overlap between $\Omega_{j,0}$ and $\Omega_{k,0}$ for $k < j$. The algorithm maintains this disjointness for $\Omega_{1,n}, \dots, \Omega_{N,n}$ for each $n \geq 1$, as well as for the limiting sets $\Omega_1, \dots, \Omega_N$. Meanwhile, the sets $\Lambda_{j,n}$ are the overlaps between $\Omega_{j,n}$ and the union of dilates $\bigcup_{m \geq 1} (A^*)^{-m} \tilde{\Omega}_n$. By translating them via T_j out to $(A^*W) \setminus W$, future overlaps should be successively smaller (as they will be compressed by $(A^*)^{-m}$ for $m \geq 1$), so that in the limit, the dilates of $\tilde{\Omega}$ by different powers of A^{*m} are disjoint. Of course, the details of this verification of condition *i* are much more complicated, but that argument in [17] (2004) is not fundamentally different from the corresponding argument for \mathbb{R}^d in [20] (1999), [21] (2001).

The proof of condition *ii*, on the other hand, requires a slight deviation from the methods of [20] (1999), [21] (2001). In both settings, the proof is relatively straightforward, because each $\Omega_{j,n}$ and Ω_j is of the form $(X \setminus Y) \cup T_j Y$, where $Y \subseteq X \subseteq W$, and X is already known to be (τ, \mathcal{D}) -congruent (or, in the \mathbb{R}^d setting, simply τ -congruent) to H^\perp . In the \mathbb{R}^d setting, the τ -congruence of the new set is immediate, because the lattice elements used for translations form a group. In our setting, with no lattice, the more complicated definition of T_j is required, with both the subtraction and the addition of an element of \mathcal{D} . The resulting (τ, \mathcal{D}) -congruence of the new set again follows easily, but the reader should note that the extra step of first subtracting the old element of \mathcal{D} is crucial. Other than that slight complication, however, the proof of condition *ii* is largely similar to those in [20] (1999), [21] (2001).

2.10 Examples of Wavelet Sets for G

We now present some examples of wavelet sets. All the examples and figures here are taken from [17] (2004). See also [24] (2003) for more examples.

Example 2.5. Let G be a LCAG with compact open subgroup H , let \mathcal{D} be a choice set of coset representatives in \hat{G} for \hat{G}/H^\perp , and let A be an expansive automorphism of G .

Take $M = 0$, so that $W = H^\perp$, set $N = |A| - 1 \geq 1$, and let $\sigma_1, \dots, \sigma_N$ be the N elements of $\mathcal{D} \cap [(A^*H^\perp) \setminus H^\perp]$. For each $j = 1, \dots, N$, define $T_j(\gamma) = \gamma - \sigma'_0 + \sigma_j$, where σ'_0 denotes the unique element of $\mathcal{D} \cap H^\perp$, and define $\Omega_{j,0} = H^\perp$. Note that $\{H^\perp, T_1 H^\perp, \dots, T_N H^\perp\}$ is a set of $|A| = N + 1$ compact open sets which together tile A^*H^\perp . See Fig. 2.8 for a diagram of $\{T_j\}$ and $\Omega_{j,0}$ ($j = 1, 2, 3$) in the case that $|A| = 4$.




Fig. 2.8 The maps T_j and the sets $\Omega_{1,0} = \Omega_{2,0} = \Omega_{3,0}$ of Example 2.5, for $|A| = 4$




Fig. 2.9 The sets $\Omega_{j,m}$ ($j = 1, 2, 3, m = 1, 2$) of Example 2.5, for $|A| = 4$

As noted in [17, Sect. 5.1], applying the algorithm of Sect. 2.9 to this data gives

$$\Omega_j = \sigma_j + H^\perp \quad \text{for all } j \in \{1, \dots, N\}.$$

Indeed, because the sets $\Omega_{1,0}, \dots, \Omega_{N,0}$ all coincide, the algorithm immediately sets every $\Omega_{j,n}$, for $j \geq 2$ and $n \geq 1$, to be the final set $\Omega_j = \sigma_j + H^\perp$. Meanwhile, the more gradual evolution of $\Omega_{1,n}$ as n increases is illustrated in Fig. 2.9; ultimately, the dark shading will cover precisely the top-most region $\Omega_1 = \sigma_1 + H^\perp$.

In this case, the choice \mathcal{D} of coset representatives is ultimately irrelevant; if σ_j and σ'_j belong to the same coset of H^\perp in \widehat{G} , then $\sigma_j + H^\perp = \sigma'_j + H^\perp$. However, as the later examples should illustrate, that happy circumstance is specific to this

example, as is the fact that we can actually write down explicit formulas for the resulting wavelets. Indeed, as noted in [17, Proposition 5.1] (2004), the wavelet generators are

$$\psi_j(x) = (x, \sigma_j) \mathbf{1}_H(x),$$

for $j = 1, \dots, N$. This simple formula leads to the surprising observation that these wavelets can be considered simultaneously to be analogs of both Haar and Shannon wavelets. See [17, Sect. 5.1] (2004) and [24, Sect. 4] (2004). They had been previously discovered in the special cases of the Cantor dyadic group by Lang in [69] (1996), and of \mathbb{Q}_p by Kozyrev in [65, Theorem 2] (2002).

Example 2.6. We can also easily produce single wavelets with the algorithm of Sect. 2.9.1. Let G be a LCAG with compact open subgroup H , let \mathcal{D} be a choice of coset representatives in \widehat{G} for \widehat{G}/H^\perp , and let A be an expansive automorphism of G .

Take $M = 0$, so that $W = H^\perp$, set $N = 1$, and let σ_1 be any one of the $|A| - 1$ elements of $\mathcal{D} \cap [(A^*H^\perp) \setminus H^\perp]$. Define $T_1(\gamma) = \gamma - \sigma'_0 + \sigma_1$, where σ'_0 denotes the unique element of $\mathcal{D} \cap H^\perp$, and define $\Omega_{1,0} = H^\perp$. See Fig. 2.10 for a diagram of T_1 and $\Omega_{1,0}$ in the case that $|A| = 4$.

As noted in [17, Sect. 5.2] (2004), it is easy to check that each $\Lambda_{1,n}$ is a translation of $(A^*)^{-n}H^\perp$, which is a dilation of H^\perp of measure $\nu(\Lambda_{1,n}) = |A|^{-n}$. Thus, each step of the algorithm translates one more successively smaller translate of $(A^*)^{-n}H^\perp$ out of H^\perp and into $H^\perp + \sigma_1$. See Fig. 2.11 for illustrations of $\Omega_{1,1}$ and $\Omega_{1,2}$ in the case $|A| = 4$; it should be easy to extrapolate what $\Omega_{1,n}$ looks like for any $n \geq 1$, and ultimately, what the wavelet set Ω_1 is.

Example 2.7. We close by giving one more example to illustrate that many other wavelets can be generated by the algorithm of Sect. 2.9.1, if one is willing to use more complicated translation functions T_j .




Fig. 2.10 The map T_1 and set $\Omega_{1,0}$ of Example 2.6, for $|A_1| = 4$




Fig. 2.11 $\Omega_{1,1}$ and $\Omega_{1,2}$ of Example 2.6

Let $G = \mathbb{Q}_3$, with compact open subgroup $H = \mathbb{Z}_3$, and let A be multiplication-by- $1/3$, so that A is expansive, with $|A| = 3$. As usual, identify \widehat{G} as \mathbb{Q}_3 and H^\perp as \mathbb{Z}_3 . Let \mathcal{S} be a set of coset representatives in \widehat{G} for \widehat{G}/H^\perp including $\sigma'_0 = 0$, $\sigma_1 = 1/3$, and $\sigma_2 = 2/3$.

Take $M = 0$, so that $W = H^\perp$, set $N = 1$, and let $\Omega_{1,0} = H^\perp$. For $\gamma \in H^\perp$, define

$$T_1(\gamma) = \begin{cases} \gamma + 2/3 & \text{if } \gamma \in 1 + 3\mathbb{Z}_3, \\ \gamma + 1/3 & \text{if } \gamma \in (3\mathbb{Z}_3) \cup (2 + 3\mathbb{Z}_3), \end{cases}$$

as in Fig. 2.12. Again, our algorithm is guaranteed to produce a single wavelet, but this time, because T_1 breaks H^\perp into two pieces before translating, the wavelet set in question is more intricate than those of Examples 2.5 and 2.6. See Figs. 2.13 and 2.14 for some of the resulting sets $\Omega_{1,m}$. Note in particular the very small disk that was moved from $\Omega_{1,2}$ to $\Omega_{1,3}$. The ultimate set Ω_1 will have successively smaller disks moved from $1 + 3\mathbb{Z}_3$ to $5/3 + 3\mathbb{Z}_3$ (i.e., from the left heavily-shaded disk of Fig. 2.12 to the right one) and from $2 + 3\mathbb{Z}_3$ to $7/3 + 3\mathbb{Z}_3$ (i.e., from the lower right lightly-shaded disk of Fig. 2.12 to the upper right one). As noted in [17, Sect. 5.3] (2004), we can describe this set explicitly as

$$\begin{aligned} \Omega_1 = & \left[\mathbb{Z}_3 \setminus \bigcup_{n=1}^{\infty} ((-5/8 + 3^{2n-2} + 3^{2n-1}\mathbb{Z}_3) \cup (-7/8 + 3^{2n-1} + 3^{2n}\mathbb{Z}_3)) \right] \\ & \cup \bigcup_{m=1}^{\infty} ((-7/24 + 3^{2m-2} + 3^{2m-1}\mathbb{Z}_3) \cup (-5/24 + 3^{2m-1} + 3^{2m}\mathbb{Z}_3)). \end{aligned}$$




Fig. 2.12 The map T_1 of Example 2.7




Fig. 2.13 $\Omega_{1,0}$ and $\Omega_{1,1}$ of Example 2.7




Fig. 2.14 $\Omega_{1,2}$ and $\Omega_{1,3}$ of Example 2.7

2.11 Epilogue

We view the construction of wavelet sets as more than a sidebar in wavelet theory and its applications. There is, of course, the sheer beauty and intricacy of many wavelet sets, and the natural questions of generalization, e.g., [20] (1999), as well as what type of theory will be required for such generalization, recalling the theories of [34] (1998) and [11] (1999) in the past.

There is also a host of geometric problems to be resolved. For example, besides the connectivity questions raised by Figs. 2.2–2.4, one would like to know if there are connected wavelet sets with connected interior. Further, there are unresolved convexity questions. We know from [23] (2006) that a wavelet set $\Omega \subseteq \mathbb{R}^d$ cannot be decomposed into a union of d or fewer convex sets, and, in particular, wavelet sets cannot be convex, see Theorem 2.10. In recent work, [79] (2008), Merrill has constructed wavelet sets $\Omega \subseteq \mathbb{R}^2$ that are finite unions of 5 or more convex sets. The lower bound “5” for \mathbb{R}^2 is not necessarily sharp, and the existence of wavelet sets $\Omega \subseteq \mathbb{R}^d$, $d > 2$, which are finite unions of convex sets is not known.

Another topic of investigation is the tantalizing relation between wavelet sets and fractals, e.g., [22] (2002), see [18] (2009) for background.

Besides the purely mathematical issues of the previous paragraphs, there is the question of applicability of wavelet sets. Naturally, one might be suspicious of ever applying the wavelet sets in Figs. 2.1 and 2.5 or the even more exotic ones in [20] (1999). However, Theorem 2.9 of Sect. 2.4, which we now repeat, provides the basis for implementation.

Theorem 2.18. *For each $n \geq 0$, $\Omega_n \setminus \Lambda_n$ is a Parseval frame wavelet set, and Ω_n is a frame wavelet set with frame bounds 1 and 2.*

In fact, sets such as $\Omega_0 \setminus \Lambda_0$ or Ω_1 can be elementary, computable shapes, and so we can construct a single wavelet frame $\{\psi_{m,n}\}$, where $\hat{\psi} = \mathbf{1}_{\Omega_1 \setminus \Lambda_1}$, say, for

$L^2(\mathbb{R}^d)$, $d \gg 0$. Further, if rapid decay of the wavelet is desirable, there are existing frame preserving smoothing results, e.g., [2] (2001), [19] (2009), [10] (2006), [51] (1997), [50] (1997), [54] (1997), and research questions, see [19] (2009). Thus, single wavelet frames can be easily constructed to give computable decompositions of the elements of $L^2(\mathbb{R}^d)$, $d \gg 0$, see Remark 2.2 in Sect. 2.1 about large data sets.

Acknowledgements The first named author gratefully acknowledges support from ONR Grant N0001409103 and MURI-ARO Grant W911NF-09-1-0383. He is also especially appreciative of wonderful mathematical interaction through the years, on the Euclidean aspect of this topic, with Professors Larry Baggett, David Larson, and Kathy Merrill, and for more recent invaluable technical contributions by Dr. Christopher Shaw. The second named author gratefully acknowledges support from NSF DMS Grant 0901494.

References

1. Patrick R. Ahern and Robert I. Jewett, *Factorization of locally compact abelian groups*, Illinois J. Math. **9** (1965), 230–235.
2. Hugo A. Aimar, Ana L. Bernardis, and Osvaldo P. Gorosito, *Perturbations of the Haar wavelet by convolution*, Proc. Amer. Math. Soc. **129** (2001), no. 12, 3619–3621 (electronic).
3. Syed Twareque Ali, Jean-Pierre Antoine, and Jean-Pierre Gazeau, *Coherent States, Wavelets and their Generalizations*, Springer-Verlag, New York, 2000.
4. Mikhail V. Altaiski, *p-adic wavelet decomposition vs Fourier analysis on spheres*, Indian J. Pure Appl. Math. **28** (1997), no. 2, 197–205.
5. Mikhail V. Altaisky, *p-adic wavelet transform and quantum physics*, Tr. Mat. Inst. Steklova **245** (2004), no. Izbr. Vopr. p-adich. Mat. Fiz. i Anal., 41–46.
6. Jean-Pierre Antoine, Yebeni B. Kouagou, Dominique Lambert, and Bruno Torr sani, *An algebraic approach to discrete dilations. Application to discrete wavelet transforms*, J. Fourier Anal. Appl. **6** (2000), no. 2, 113–141.
7. Jean-Pierre Antoine and Pierre Vandergheynst, *Wavelets on the 2-sphere: a group-theoretical approach*, Appl. Comput. Harmon. Anal. **7** (1999), no. 3, 262–291.
8. Pascal Auscher, *Solution of two problems on wavelets*, J. Geom. Anal. **5** (1995), no. 2, 181–236.
9. Larry Baggett, Alan Carey, William Moran, and Peter Ohring, *General existence theorems for orthonormal wavelets, an abstract approach*, Publ. Res. Inst. Math. Sci. **31** (1995), no. 1, 95–111.
10. Lawrence W. Baggett, Palle E. T. Jorgensen, Kathy D. Merrill, and Judith A. Packer, *Construction of Parseval wavelets from redundant filter systems*, J. Math. Phys. **46** (2005), no. 8, 083502, 28.
11. Lawrence W. Baggett, Herbert A. Medina, and Kathy D. Merrill, *Generalized multi-resolution analyses and a construction procedure for all wavelet sets in \mathbb{R}^n* , J. Fourier Anal. Appl. **5** (1999), no. 6, 563–573.
12. Lawrence W. Baggett and Kathy D. Merrill, *Abstract harmonic analysis and wavelets in \mathbb{R}^n* , The functional and harmonic analysis of wavelets and frames (San Antonio, TX, 1999), Amer. Math. Soc., Providence, RI, 1999, pp. 17–27.
13. John J. Benedetto, *Idele characters in spectral synthesis on $\mathbb{R}/2\pi\mathbb{Z}$* , Ann. Inst. Fourier **23** (1973), 43–64.
14. ———, *Zeta functions for idelic pseudo-measures*, Ann. Scuola Norm. Sup. Pisa Cl. Sci. **4** (1979), no. 6, 367–377.
15. ———, *Harmonic Analysis and Applications*, Studies in Advanced Mathematics, CRC Press, Boca Raton, FL, 1997.
16. ———, *Frames, sampling, and seizure prediction*, Advances in Wavelets, edited by K.-S. Lau, Springer-Verlag, New York, 1998.

17. John J. Benedetto and Robert L. Benedetto, *A wavelet theory for local fields and related groups*, J. Geom. Anal. **14** (2004), 423–456.
18. John J. Benedetto and Wojciech Czaja, *Integration and Modern Analysis*, Birkhäuser Advanced texts, Birkhäuser Boston, MA, 2009.
19. John J. Benedetto and Emily J. King, *Smooth functions associated with wavelet sets on \mathbb{R}^d , $d \geq 1$, and frame bound gaps*, Acta Appl. Math. **107** (2009), no. 1-3, 121–142.
20. John J. Benedetto and Manuel T. Leon, *The construction of multiple dyadic minimally supported frequency wavelets on \mathbb{R}^d* , Contemporary Mathematics **247** (1999), 43–74.
21. ———, *The construction of single wavelets in d -dimensions*, J. of Geometric Analysis **11** (2001), 1–15.
22. John J. Benedetto and Songkiat Sumetkijakan, *A fractal set constructed from a class of wavelet sets*, Inverse problems, image analysis, and medical imaging (New Orleans, LA, 2001), Contemp. Math., vol. 313, Amer. Math. Soc., Providence, RI, 2002, pp. 19–35.
23. ———, *Tight frames and geometric properties of wavelet sets*, Adv. Comput. Math. **24** (2006), no. 1-4, 35–56.
24. Robert L. Benedetto, *Examples of wavelets for local fields*, Wavelets, Frames, and Operator Theory (College Park, MD, 2003), Amer. Math. Soc., Providence, RI, 2004, pp. 27–47.
25. Andrea Calogero, *A characterization of wavelets on general lattices*, J. Geom. Anal. **10** (2000), no. 4, 597–622.
26. Ole Christensen, *Frames, Riesz bases, and discrete Gabor/wavelet expansions*, Bull. Amer. Math. Soc. (N.S.) **38** (2001), no. 3, 273–291 (electronic).
27. Charles K. Chui, Wojciech Czaja, Mauro Maggioni, and Guido Weiss, *Characterization of general tight frames with matrix dilations and tightness preserving oversampling*, J. Fourier Anal. Appl. **8** (2002), no. 2, 173–200.
28. Albert Cohen and Ingrid Daubechies, *Nonseparable bidimensional wavelet bases*, Rev. Mat. Iberoamericana **9** (1993), no. 1, 51–137.
29. Ronald R. Coifman and Guido Weiss, *Book Review: Littlewood-Paley and Multiplier Theory*, Bull. Amer. Math. Soc. **84** (1978), no. 2, 242–250.
30. Jennifer Courter, *Construction of dilation- d wavelets*, AMS Contemporary Mathematics **247** (1999), 183–205.
31. Stephan Dahlke, *Multiresolution analysis and wavelets on locally compact abelian groups*, Wavelets, images, and surface fitting (Chamonix-Mont-Blanc, 1993), A K Peters, Wellesley, MA, 1994, pp. 141–156.
32. ———, *The construction of wavelets on groups and manifolds*, General algebra and discrete mathematics (Potsdam, 1993), Heldermann, Lemgo, 1995, pp. 47–58.
33. Xingde Dai, Yuanan Diao, Qing Gu, and Deguang Han, *Frame wavelets in subspaces of $L^2(\mathbb{R}^d)$* , Proc. Amer. Math. Soc. **130** (2002), no. 11, 3259–3267 (electronic).
34. Xingde Dai and David R. Larson, *Wandering vectors for unitary systems and orthogonal wavelets*, Mem. Amer. Math. Soc. **134** (1998), no. 640, viii+68.
35. Xingde Dai, David R. Larson, and Darrin M. Speegle, *Wavelet sets in \mathbb{R}^n* , J. Fourier Anal. Appl. **3** (1997), no. 4, 451–456.
36. ———, *Wavelet sets in \mathbb{R}^n . II*, Wavelets, multiwavelets, and their applications (San Diego, CA, 1997), Contemp. Math., vol. 216, Amer. Math. Soc., Providence, RI, 1998, pp. 15–40.
37. James Daly and Keith Phillips, *On singular integrals, multipliers, \mathfrak{H}^1 and Fourier series — a local field phenomenon*, Math. Ann. **265** (1983), no. 2, 181–219.
38. Ingrid Daubechies, *Ten Lectures on Wavelets*, CBMS-NSF Series in Applied Math. SIAM, Philadelphia, PA, 1992.
39. Ingrid Daubechies and Bin Han, *The canonical dual frame of a wavelet frame*, Appl. Comput. Harmon. Anal. **12** (2002), no. 3, 269–285.
40. Richard J. Duffin and Albert C. Schaeffer, *A class of nonharmonic Fourier series*, Trans. Amer. Math. Soc. **72** (1952), 341–366.
41. Robert E. Edwards and Garth I. Gaudry, *Littlewood-Paley and multiplier theory*, Springer-Verlag, Berlin, 1977, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 90.
42. Xiang Fang and Xihua Wang, *Construction of minimally supported frequency wavelets*, J. Fourier Anal. Appl. **2** (1996), no. 4, 315–327.

43. Yuri A. Farkov, *Orthogonal wavelets on locally compact abelian groups*, Funktsional. Anal. i Prilozhen. **31** (1997), no. 4, 86–88.
44. ———, *Orthogonal p -wavelets on \mathbb{R}_+* , Wavelets and Splines, St. Petersburg Univ. Press, St. Petersburg, 2005, pp. 4–26.
45. ———, *Orthogonal wavelets with compact supports on locally compact abelian groups*, Izv. Ross. Akad. Nauk Ser. Mat. **69** (2005), no. 3, 193–220.
46. Michael Frazier, Björn Jawerth, and Guido Weiss, *Littlewood-Paley Theory and the Study of Function Spaces*, CBMS Regional Conference Series in Mathematics, vol. 79, Amer. Math. Soc., Washington, DC, 1991.
47. Gustaf Gripenberg, *A necessary and sufficient condition for the existence of a father wavelet*, Studia Math. **114** (1995), no. 3, 207–226.
48. Karlheinz Gröchenig and Wolodymyr R. Madych, *Multiresolution analysis, Haar bases, and self-similar tilings of $L^2(\mathbb{R}^n)$* , IEEE Trans. Information Theory **38**(2) (1992), 556–568.
49. Qing Gu and Deguang Han, *On multiresolution analysis (MRA) wavelets in \mathbb{R}^n* , J. Fourier Anal. Appl. **6** (2000), no. 4, 437–447.
50. Bin Han, *On dual wavelet tight frames*, Appl. Comput. Harmon. Anal. **4** (1997), no. 4, 380–413.
51. Deguang Han, *Unitary systems, wavelets, and operator algebras*, Ph.D. thesis, Texas A&M University, College Station, TX, 1997.
52. Deguang Han, David R. Larson, Manos Papadakis, and Theodoros Stavropoulos, *Multiresolution analyses of abstract Hilbert spaces and wandering subspaces*, The functional and harmonic analysis of wavelets and frames (San Antonio, TX, 1999), Amer. Math. Soc., Providence, RI, 1999, pp. 259–284.
53. Eugenio Hernández, Xihua Wang, and Guido Weiss, *Smoothing minimally supported frequency wavelets. I*, J. Fourier Anal. Appl. **2** (1996), no. 4, 329–340.
54. ———, *Smoothing minimally supported frequency wavelets, part II*, J. Fourier Anal. Appl. **3** (1997), 23–41.
55. Eugenio Hernández and Guido Weiss, *A First Course on Wavelets*, Studies in Advanced Mathematics, CRC Press, Boca Raton, FL, 1996, With a foreword by Yves Meyer.
56. Edwin Hewitt and Kenneth A. Ross, *Abstract Harmonic Analysis, Volume I*, Springer-Verlag, New York, 1963.
57. ———, *Abstract Harmonic Analysis, Volume II*, Springer-Verlag, New York, 1970.
58. Matthias Holschneider, *Wavelet analysis over abelian groups*, Appl. Comput. Harmon. Anal. **2** (1995), no. 1, 52–60.
59. Huikun Jiang, Dengfeng Li, and Ning Jin, *Multiresolution analysis on local fields*, J. Math. Anal. Appl. **294** (2004), no. 2, 523–532.
60. Carolyn P. Johnston, *On the pseudodilation representations of Flornes, Grossmann, Holschneider, and Torresani*, J. Fourier Anal. Appl. **3** (1997), no. 4, 377–385.
61. Andrei Yu. Khrennikov and Sergei V. Kozyrev, *Wavelets on ultrametric spaces*, Appl. Comput. Harmon. Anal. **19** (2005), no. 1, 61–76.
62. Andrei Yu. Khrennikov, Vladimir M. Shelkovich, and Maria Skopina, *p -adic refinable functions and MRA-based wavelets*, J. Approx. Theory **161** (2009), no. 1, 226–238.
63. Andrei Kolmogorov, *Une contribution à l'étude de la convergence des séries de Fourier*, Fund. Math. **5** (1924), 96–97.
64. Jelena Kovačević and Martin Vetterli, *Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for \mathbb{R}^n* , IEEE Trans. Inform. Theory **38** (1992), no. 2, part 2, 533–555.
65. Sergei V. Kozyrev, *Wavelet analysis as a p -adic spectral analysis*, Izv. Math **66** (2002), 367–376.
66. ———, *p -adic pseudodifferential operators and p -adic wavelets*, Theoret. and Math. Phys. **138** (2004), no. 3, 322–332.
67. Jeffrey C. Lagarias and Yang Wang, *Integral self-affine tiles in \mathbb{R}^n , part II lattice tilings*, J. Fourier Anal. Appl. **3** (1997), 83–102.
68. Jeffrey C. Lagarias and Yang Wang, *Orthogonality criteria for compactly supported refinable functions and refinable function vectors*, J. Fourier Anal. Appl. **6** (2000), no. 2, 153–170.

69. W. Christopher Lang, *Orthogonal wavelets on the Cantor dyadic group*, SIAM J. Math. Anal. **27** (1996), 305–312.
70. ———, *Fractal multiwavelets related to the Cantor dyadic group*, Internat. J. Math. Math. Sci. **21** (1998), 307–314.
71. ———, *Wavelet analysis on the Cantor dyadic group*, Houston J. Math. **24** (1998), 533–544 and 757–758.
72. Pierre G. Lemarié, *Base d'ondelettes sur les groupes de Lie stratifiés*, Bull. Soc. Math. France **117** (1989), no. 2, 211–232.
73. Yunzhang Li, *On the holes of a class of bidimensional nonseparable wavelets*, J. Approx. Theory **125** (2003), no. 2, 151–168.
74. John E. Littlewood and Raymond E. A. C. Paley, *Theorems on Fourier series and power series*, J. London Math. Soc. **6** (1931), 230–233.
75. ———, *Theorems on Fourier series and power series III*, Proc. London Math. Soc. **43** (1937), 105–126.
76. Taras P. Lukashenko, *Wavelets on topological groups*, Izv. Ross. Akad. Nauk Ser. Mat. **58** (1994), no. 3, 88–102.
77. Wolodymyr R. Madych, *Some elementary properties of multiresolution analyses of $L^2(\mathbb{R}^n)$* , Wavelets: a Tutorial in Theory and Applications, edited by Charles K. Chui, Academic Press, Inc, San Diego, CA, 1992, pp. 256–294.
78. Stéphane G. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, Boston, 1998.
79. Kathy D. Merrill, *Simple wavelet sets for integral dilations in \mathbb{R}^2* , Representations, Wavelets, and Frames: A Celebration of the Mathematical Work of Lawrence W. Baggett, Birkhäuser, Boston, MA, 2008, pp. 177–192.
80. Yves Meyer, *Ondelettes, fonctions splines et analyses graduées*, Rend. Sem. Mat. Univ. Politec. Torino **45** (1987), no. 1, 1–42 (1988).
81. ———, *Ondelettes et Opérateurs*, Hermann, Paris, 1990.
82. ———, *Wavelets and Operators*, Cambridge Studies in Advanced Mathematics, vol. 37, Cambridge University Press, Cambridge, 1992, Translated from the 1990 French original by D. H. Salinger.
83. Cui Minggen, GuangHong Gao, and Phil Ung Chung, *On the wavelet transform in the field \mathbb{Q}_p of p -adic numbers*, Appl. Comput. Harmon. Anal. **13** (2002), no. 2, 162–168.
84. Manos Papadakis, *Generalized frame multiresolution analysis of abstract Hilbert spaces*, Sampling, Wavelets, and Tomography, Birkhäuser, Boston, MA, 2003.
85. Manos Papadakis and Theodoros Stavropoulos, *On the multiresolution analyses of abstract Hilbert spaces*, Bull. Greek Math. Soc. **40** (1998), 79–92.
86. Lev Lemenovich Pontryagin, *Topological Groups, 2nd edition*, Gordon and Breach, Science Publishers, Inc., New York, 1966, Translated from the Russian by Arlen Brown.
87. Hans Reiter, *Classical Harmonic Analysis and Locally Compact Groups*, Oxford University Press, 1968.
88. Walter Rudin, *Fourier Analysis on Groups*, John Wiley and Sons, New York, 1962.
89. Eckart Schulz and Keith F. Taylor, *Extensions of the Heisenberg group and wavelet analysis in the plane*, Spline functions and the theory of wavelets (Montreal, PQ, 1996), Amer. Math. Soc., Providence, RI, 1999, pp. 217–225. MR 99m:42053
90. Vladimir Shelkovich and Maria Skopina, *p -adic Haar multiresolution analysis and pseudo-differential operators*, J. Fourier Anal. Appl. **15** (2009), no. 3, 366–393.
91. Paolo M. Soardi and David Weiland, *Single wavelets in n -dimensions*, J. Fourier Anal. Appl. **4** (1998), 299–315.
92. Elias M. Stein, *Singular Integrals and Differentiability Properties of Functions*, Princeton Mathematical Series, No. 30, Princeton University Press, Princeton, N.J., 1970.
93. ———, *Topics in Harmonic Analysis Related to the Littlewood-Paley Theory.*, Annals of Mathematics Studies, No. 63, Princeton University Press, Princeton, N.J., 1970.
94. Elias M. Stein and Guido Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*, Princeton University Press, Princeton, N.J., 1971, Princeton Mathematical Series, No. 32.
95. Robert S. Strichartz, *Construction of orthonormal wavelets*, Wavelets: Mathematics and Applications, Stud. Adv. Math., CRC, Boca Raton, FL, 1994, pp. 23–50.

96. Mitchell H. Taibleson, *Fourier Analysis on Local Fields*, Princeton U. Press, Princeton, NJ, 1975.
97. Khalifa Trimèche, *Continuous wavelet transform on semisimple Lie groups and inversion of the Abel transform and its dual*, Collect. Math. **47** (1996), no. 3, 231–268.
98. Victor G. Zakharov, *Nonseparable multidimensional Littlewood-Paley like wavelet bases*, Preprint, 1996.

Chapter 3

The Measure of the Closure of a Wavelet Set May Be $> 2\pi$

Zhihua Zhang

Abstract In this paper, we will show that the Fourier support of a wavelet ψ possesses the minimal measure if and only if the conditions $|\widehat{\psi}| = \chi_E$ and $|E| = |\overline{E}| = 2\pi$ hold simultaneously. If a wavelet ψ satisfies the first condition $|\widehat{\psi}| = \chi_E$, then we call E a wavelet set. In order to explain why the second condition is indispensable, we will construct a wavelet set E^* such that the measure of its closure satisfies $|\overline{E^*}| > |E^*|$.

3.1 Introduction

It is well known that the minimal measure of the Fourier supports of wavelets is 2π . In fact, if ψ is a wavelet, then $\sum_k |\widehat{\psi}(\omega + 2k\pi)|^2 = 1$ and so $|\widehat{\psi}(\omega)| \leq 1$. Furthermore, we have

$$2\pi = \int_{\mathbf{R}} |\widehat{\psi}(t)|^2 dt = \int_E |\widehat{\psi}(t)|^2 dt \leq |E| \leq |\overline{E}| = |\text{supp } \widehat{\psi}|, \quad (3.1)$$

where $E = \{\omega \in \mathbf{R} \mid \widehat{\psi}(\omega) \neq 0\}$ and \overline{E} is the closure of E . From this, we see that the Fourier support of a wavelet ψ possesses the minimal measure if and only if

$$|\widehat{\psi}| = \chi_E \quad \text{and} \quad |E| = |\overline{E}| = 2\pi,$$

where χ_E is the characteristic function of E . If a wavelet ψ satisfies the first condition $|\widehat{\psi}| = \chi_E$, then we call E a wavelet set. Many wavelet sets have been constructed

Z. Zhang (✉)

College of Global Change and Earth System Science, Beijing Normal University,
Beijing 100875, China

and

Department of Mathematics, University of California, Davis, CA 95616, USA

e-mail: zhangzh@math.ucdavis.edu

in [1–5, 7]. For example, the following set

$$E = \left[-\frac{32\pi}{7}, -4\pi\right) \cup \left[-\pi, -\frac{4\pi}{7}\right) \cup \left[\frac{4\pi}{7}, \pi\right) \cup \left[4\pi, \frac{32\pi}{7}\right)$$

is a wavelet set. It is obvious that if a wavelet set E is a union of finitely many intervals, then the wavelet set E must satisfy the second condition $|E| = |\bar{E}| = 2\pi$. Moreover, all of known wavelet sets, which consist of infinitely many intervals, also satisfy $|E| = |\bar{E}| = 2\pi$. Hence, it is natural to ask whether there exists a wavelet set E^* such that $|\bar{E}^*| > |E^*| = 2\pi$. In this paper, we will give a positive answer in Theorem 3.1.

For a scaling function, there exists a similar problem. We can prove that the Fourier support of a scaling function φ possesses the minimal measure if and only if $|\hat{\varphi}| = \chi_G$ and $|G| = |\bar{G}|$. For the same reason, it is also natural to ask whether there exists a scaling function φ such that

$$|\hat{\varphi}| = \chi_G \quad \text{and} \quad |\text{supp } \hat{\varphi}| = |\bar{G}| > |G| = 2\pi.$$

In Theorem 3.2, we will answer this problem.

3.2 Preliminaries

Define the Fourier transform of $f \in L^2(\mathbf{R})$ by \hat{f} :

$$\hat{f}(\omega) = \int_{\mathbf{R}} f(t) e^{-it\omega} dt.$$

For convenience, the support of the Fourier transform of f is called the Fourier support of f , i.e.,

$$\text{supp } \hat{f} = \text{clos} \{ \omega \in \mathbf{R} \mid \hat{f}(\omega) \neq 0 \}.$$

A function $\psi \in L^2(\mathbf{R})$ is called a wavelet if $\{\psi_{j,k} : j, k \in \mathbf{Z}\}$ is an orthonormal basis for $L^2(\mathbf{R})$, where $\psi_{j,k} = 2^{\frac{j}{2}} \psi(2^j \cdot -k)$.

A multiresolution analysis (MRA) [3] consists of a sequence of subspaces $\{V_m\}_m$ of $L^2(\mathbf{R})$ satisfying

- (1) $V_m \subset V_{m+1}$, $\bigcup_m V_m = L^2(\mathbf{R})$, $\bigcap_m V_m = \{0\}$
- (2) $f \in V_m \iff f(2 \cdot) \in V_{m+1}$, $m \in \mathbf{Z}$
- (3) there exists a $\varphi \in V_0$ such that $\{\varphi(\cdot - n)\}_n$ is an orthonormal basis of V_0 , where φ is called the scaling function

If a wavelet ψ belongs to the space W_0 , where $V_0 \oplus W_0 = V_1$, then ψ is called a MRA wavelet.

Let E be a set of \mathbf{R} . Denote

$$E + b = \{ \omega + b, \omega \in E \} \quad (b \in \mathbf{R}) \quad \text{and} \quad cE = \{ c\omega, \omega \in E \} \quad (c \in \mathbf{R})$$

Let E_1 and E_2 be two sets. We say $E_1 \simeq E_2$ if the measure

$$|E_1 \setminus E_2| = |E_2 \setminus E_1| = 0.$$

A set E is called a wavelet set if there exists a wavelet ψ such that $|\hat{\psi}| = \chi_E$, where χ_E is the characteristic function of E .

We need the following lemmas.

Lemma 3.1. [4] *A set E is a wavelet set if and only if $\{E + 2k\pi, k \in \mathbf{Z}\}$ is a partition of \mathbf{R} and $\{2^j E, j \in \mathbf{Z}\}$ is a partition of \mathbf{R} .*

Lemma 3.2. [5] *Let ψ be a wavelet derived by the scaling function ϕ . Then*

$$|\hat{\phi}(\omega)|^2 = \sum_{m \geq 1} |\hat{\psi}(2^m \omega)|^2 \quad a.e. \ \omega \in \mathbf{R}.$$

Lemma 3.3. [2] *Let ψ be a wavelet with $\text{supp } \hat{\psi} \subset [-\frac{8\alpha}{3}, 4\pi - \frac{4\alpha}{3}]$ ($0 < \alpha \leq \pi$). Then ψ is a MRA wavelet.*

3.3 A Wavelet Set E with $|\bar{E}| > 2\pi$

In this section, we will construct a wavelet set E such that $|\bar{E}| > |E| = 2\pi$.

Theorem 3.1. *There exists a wavelet set E with $|\bar{E}| > |E| = 2\pi$.*

Proof. At first, we construct a set E . Take a countable, dense set $\{\omega_n\}_1^\infty$ in the interval $(\frac{2\pi}{3}, \frac{4\pi}{3})$. Denote the point sets $\{F_i\}_1^4$ as follows.

$$F_1 = \bigcup_{n=1}^\infty \left(\left(\omega_n - \frac{\pi}{2^{n+2}}, \omega_n + \frac{\pi}{2^{n+2}} \right) \cap \left(\frac{2\pi}{3}, \frac{4\pi}{3} \right) \right) \quad (3.2)$$

and

$$F_2 = \left[\frac{4\pi}{3}, \frac{8\pi}{3} \right] \setminus 2F_1, \quad F_3 = \left[-\frac{4\pi}{3}, -\frac{2\pi}{3} \right] \setminus (F_1 - 2\pi), \quad F_4 = 2F_1 - 4\pi. \quad (3.3)$$

Define $E = \bigcup_{i=1}^4 F_i$. It is easy to check that

$$\begin{aligned} F_1 &\subset \left(\frac{2\pi}{3}, \frac{4\pi}{3} \right), & F_2 &\subset \left[\frac{4\pi}{3}, \frac{8\pi}{3} \right], \\ F_3 &\subset \left[-\frac{4\pi}{3}, -\frac{2\pi}{3} \right], & F_4 &\subset \left(-\frac{8\pi}{3}, -\frac{4\pi}{3} \right) \end{aligned}$$

and

$$|F_2| = \frac{4\pi}{3} - 2|F_1|, \quad |F_3| = \frac{2\pi}{3} - |F_1|, \quad |F_4| = 2|F_1|.$$

From this, we see that the sets $\{F_i\}_1^4$ are pairwise disjoint and

$$E \subset \left(-\frac{8\pi}{3}, -\frac{2\pi}{3}\right] \cup \left(\frac{2\pi}{3}, \frac{8\pi}{3}\right], \quad |E| = \sum_{i=1}^4 |F_i| = 2\pi. \quad (3.4)$$

(1) We prove that $\{2^m E, m \in \mathbf{Z}\}$ is a partition of \mathbf{R} .

Since

$$\left(F_1 \cup \frac{1}{2}F_2\right) \cup \left(2F_1 \cup F_2\right) = \left[\frac{2\pi}{3}, \frac{8\pi}{3}\right]$$

and

$$\left(F_3 \cup \frac{1}{2}F_4\right) \cup \left(2F_3 \cup F_4\right) = \left[-\frac{8\pi}{3}, -\frac{2\pi}{3}\right],$$

we have

$$\mathbf{R} \supset \bigcup_m 2^m E \supset \left(\bigcup_l 2^l \left[-\frac{8\pi}{3}, -\frac{2\pi}{3}\right]\right) \cup \left(\bigcup_l 2^l \left[\frac{2\pi}{3}, \frac{8\pi}{3}\right]\right) = \mathbf{R} \setminus \{0\}.$$

So

$$\bigcup_m 2^m E \simeq \mathbf{R}. \quad (3.5)$$

Below we prove $(2^m E) \cap (2^{m'} E) \simeq \emptyset$ ($m \neq m'$, $m, m' \in \mathbf{Z}$). Without loss of generality, we only prove

$$(2^m E) \cap E \simeq \emptyset \quad (m \neq 0).$$

We first prove $(2^m F_3) \cap E \simeq \emptyset$ ($m \neq 0$).

For $m \leq -1$, $(2^m F_3) \subset [-\frac{2\pi}{3}, 0]$, so

$$(2^m F_3) \cap E \simeq \emptyset \quad (m \leq -1). \quad (3.6)$$

For $m \geq 1$, we have $(2^m F_3) \subset (-\infty, -\frac{4\pi}{3}]$, so

$$(2^m F_3) \cap F_i \simeq \emptyset \quad (i \neq 4).$$

In case $i = 4$, we have

$$(2^m F_3 \cap F_4) \subset \left(-\infty, -\frac{8\pi}{3}\right] \cap \left(-\frac{8\pi}{3}, -\frac{4\pi}{3}\right) = \emptyset \quad (m \geq 2)$$

and

$$2F_3 \cap F_4 = 2F_3 \cap (2F_1 - 4\pi) = 2(F_3 \cap (F_1 - 2\pi)) = \emptyset,$$

so

$$(2^m F_3) \cap F_4 = \emptyset \quad (m \geq 1).$$

Finally, we get

$$(2^m F_3) \cap E \simeq \emptyset \quad (m \geq 1).$$

Again by (3.6), we get $(2^m F_3) \cap E \simeq \emptyset$ ($m \neq 0$).

Similarly, we have

$$(2^m F_i) \cap E \simeq \emptyset \quad (m \neq 0, i = 1, 2, 4),$$

and so

$$(2^m E) \cap E \simeq \emptyset \quad (m \neq 0).$$

Therefore, $\{2^m E, m \in \mathbf{Z}\}$ is a partition of \mathbf{R} (a.e.).

(2) We prove that $\{E + 2m\pi, m \in \mathbf{Z}\}$ is a partition of \mathbf{R} .

Since

$$(F_1 \cup (F_3 + 2\pi)) \cup (F_2 \cup (F_4 + 4\pi)) = \left[\frac{2\pi}{3}, \frac{8\pi}{3} \right],$$

we have

$$\bigcup_m (E + 2m\pi) = \bigcup_m \left(\left[\frac{2\pi}{3}, \frac{8\pi}{3} \right] + 2m\pi \right) = \mathbf{R}.$$

Below we prove $(E + 2m\pi) \cap (E + 2m'\pi) = \emptyset$ ($m \neq m', m, m' \in \mathbf{Z}$). Without loss of generality, we only prove

$$(E + 2m\pi) \cap E = \emptyset \quad (m \neq 0).$$

We first prove that $(F_1 + 2m\pi) \cap E = \emptyset$ ($m \neq 0$).

For $m \geq 1$, since $F_1 + 2m\pi \subset (\frac{8\pi}{3}, \infty)$,

$$(F_1 + 2m\pi) \cap E = \emptyset \quad (m \geq 1). \quad (3.7)$$

For $m \leq -2$, since $F_1 + 2m\pi \subset (-\infty, -\frac{8\pi}{3})$, we have

$$(F_1 + 2m\pi) \cap E = \emptyset \quad (m \leq -2). \quad (3.8)$$

From $(F_1 - 2\pi) \cap F_3 = \emptyset$ and

$$((F_1 - 2\pi) \cap F_i) \subset \left(\left(-\frac{4\pi}{3}, -\frac{2\pi}{3} \right) \cap F_i \right) = \emptyset \quad (i = 1, 2, 4),$$

it follows that

$$(F_1 - 2\pi) \cap E = \emptyset.$$

Again by (3.8) and (3.7), we get $(F_1 + 2m\pi) \cap E = \emptyset$ ($m \neq 0$).

Similarly, we have $(F_i + 2m\pi) \cap E = \emptyset$ ($m \neq 0, i = 2, 3, 4$), and so

$$(E + 2m\pi) \cap E = \emptyset \quad (m \neq 0).$$

Therefore, $\{E + 2m\pi, m \in \mathbf{Z}\}$ is a partition of \mathbf{R} .

Using Lemma 3.1, we know that E is a wavelet set.

(3) We prove that $|\overline{E}| > 2\pi$. By (3.2) and (3.3), we have

$$\begin{aligned} \overline{F}_1 &= \left[\frac{2\pi}{3}, \frac{4\pi}{3} \right], & \overline{F}_2 &\subset \left[\frac{4\pi}{3}, \frac{8\pi}{3} \right], \\ \overline{F}_3 &\subset \left[-\frac{4\pi}{3}, -\frac{2\pi}{3} \right], & \overline{F}_4 &= 2\overline{F}_1 - 4\pi = \left[-\frac{8\pi}{3}, -\frac{4\pi}{3} \right]. \end{aligned}$$

From this, we see that the four sets $\{\overline{F}_i\}$ are pairwise disjoint sets except the two points $-\frac{4\pi}{3}$ and $\frac{4\pi}{3}$. Again by $E = \bigcup_{i=1}^4 F_i$, we get

$$|\overline{E}| = \left| \bigcup_{i=1}^4 \overline{F}_i \right| = \sum_{i=1}^4 |\overline{F}_i|. \quad (3.9)$$

Since $\{\omega_n\}_1^\infty \subset F_1$ and $\{\omega_n\}_1^\infty$ is dense in $(\frac{2\pi}{3}, \frac{4\pi}{3})$, we have $\overline{F}_1 = [\frac{2\pi}{3}, \frac{4\pi}{3}]$, so $|\overline{F}_1| = \frac{2\pi}{3}$. Again by (3.2), we have $|F_1| \leq \frac{\pi}{2}$. So

$$|\overline{F}_1| > |F_1|. \quad (3.10)$$

Again by (3.4), (3.9), (3.10), and $|\overline{F}_i| \geq |F_i|$ ($i = 2, 3, 4$), we get

$$|\overline{E}| = \sum_{i=1}^4 |\overline{F}_i| > \sum_{i=1}^4 |F_i| = |E| = 2\pi. \quad (3.11)$$

Theorem 3.1 is proved. \square

From Theorem 3.1 and (3.1), we see that in the characterization “the measure of the Fourier support of a wavelet ψ is minimal if and only if $\widehat{\psi} = \chi_F$ and $|\overline{F}| = |F|$ ”, the condition $|\overline{F}| = |F|$ is indispensable.

3.4 Fourier Supports of Scaling Functions

For scaling functions, there exists a similar problem. Let φ be a scaling function. Then we have

$$2\pi = \int_{\mathbf{R}} |\widehat{\varphi}(\omega)|^2 d\omega = \int_G |\widehat{\varphi}(\omega)|^2 d\omega \leq |G| \leq |\overline{G}| = |\text{supp } \widehat{\varphi}|.$$

From this, we see that the Fourier support of a scaling function φ possesses the minimal measure if and only if $|\widehat{\varphi}| = \chi_G$ and $|\overline{G}| = |G|$. The scaling function φ satisfying the first condition $|\widehat{\varphi}| = \chi_G$ have been researched by Madych and the MRA associated with it is called a translation invariant MRA [6].

In the following theorem, we claim that the second condition is indispensable.

Theorem 3.2. *There exists a scaling function φ with $\widehat{\varphi} = \chi_G$ such that $|\overline{G}| > |G| = 2\pi$.*

Proof. Let E be a wavelet set which is stated in Theorem 3.1. Define ψ by $\widehat{\psi} = \chi_E$. Then ψ is a wavelet and $\text{supp } \widehat{\psi} = E$. Since $E \subset [-\frac{8\pi}{3}, \frac{8\pi}{3}]$, by Lemma 3.3, we know that ψ is an MRA wavelet.

Now we give the corresponding scaling function φ . By Lemma 3.2 and $\widehat{\psi} = \chi_E$,

$$|\widehat{\varphi}(\omega)|^2 = \sum_{m=1}^{\infty} \chi_E(2^m \omega) \quad a.e. \omega \in \mathbf{R}. \quad (3.12)$$

since $\{2^{-m}E\}_1^{\infty}$ are pairwise disjoint. By (3.12), we can take $\widehat{\varphi}(\omega) = \chi_G(\omega)$, where

$$G = \bigcup_{m=1}^{\infty} 2^{-m}E \quad \text{and} \quad E = \bigcup_{i=1}^4 F_i. \quad (3.13)$$

where $\{F_i\}_1^4$ are defined in Sect. 3.3.

By (3.2) and (3.3), we have $F_1 \cup \frac{1}{2}F_2 = [\frac{2\pi}{3}, \frac{4\pi}{3}]$, so

$$\frac{1}{2^m}F_1 \cup \frac{1}{2^{m+1}}F_2 = \left[\frac{2\pi}{3 \cdot 2^m}, \frac{4\pi}{3 \cdot 2^m} \right]$$

and

$$\bigcup_{m=0}^{\infty} \left(\frac{1}{2^m}F_1 \cup \frac{1}{2^{m+1}}F_2 \right) = \left(0, \frac{4\pi}{3} \right].$$

This implies that

$$\bigcup_{m=1}^{\infty} \frac{1}{2^m} (F_1 \cup F_2) = \left(0, \frac{4\pi}{3} \right] \setminus F_1.$$

By $F_3 \cup \frac{1}{2}F_4 = [-\frac{4\pi}{3}, -\frac{2\pi}{3}]$, we have

$$\bigcup_{m=1}^{\infty} \frac{1}{2^m} (F_3 \cup F_4) = \left[-\frac{4\pi}{3}, 0\right) \setminus F_3 = \left[-\frac{2\pi}{3}, 0\right) \cup (F_1 - 2\pi).$$

From this and (3.13), we have

$$G = \bigcup_{m=1}^{\infty} 2^{-m} (F_1 \cup F_2 \cup F_3 \cup F_4) = (F_1 - 2\pi) \cup P,$$

where $P = \left([- \frac{2\pi}{3}, 0) \cup (0, \frac{4\pi}{3}]\right) \setminus F_1$. So the scaling function ϕ with

$$\overline{G} = \text{supp } \widehat{\phi} = (\overline{F}_1 - 2\pi) \cup \overline{P}.$$

However, by (3.10) and $|\overline{P}| \geq 2\pi - |F_1|$, we have

$$|\overline{G}| = |\text{supp } \widehat{\phi}| = |\overline{F}_1 - 2\pi| + |\overline{P}| = |\overline{F}_1| + |\overline{P}| > |F_1| + 2\pi - |F_1| = 2\pi.$$

Theorem 3.2 is proved. \square

References

1. Hernández, E., Wang, X., and Weiss G., Smoothing minimally supported frequency wavelets, Part 1, J. Fourier Anal. and Appl., **2**, 329–340(1996)
2. Brandolini, L., Garrigos, G., Rzeszotnik, Z. and Weiss, G.: The behavior at the origin of a class of band-limited wavelets, Contemp. Math. **247**, 75–91(1999)
3. Daubechies, I.: *Ten Lectures on Wavelets*, CBMS- Conference Lecture Notes, SIAM, Philadelphia (1992)
4. Hernández, E. and Weiss, G.: *A First Course on Wavelets*, CRC Press, New York (1996).
5. Long, R.: *Multidimensional Wavelet Analysis*, International publishing Co. Beijing (1995)
6. Madych, W.: Some elementary properties of multiresolution analyses of $L^2(\mathbf{R}^n)$. In Chui, C. K. (ed.) *Wavelets: A Tutorial in Theory and Applications*, pp. 259–294, Academic Press, New York (1992).
7. Speegle, D. M.: The S -elementary wavelets are path-connected, Proc. Amer. Math. Soc., **127**, 223–233(1999).

Chapter 4

Quincunx Wavelets on \mathbb{T}^2

Kenneth R. Hoover and Brody Dylan Johnson

Abstract This article examines a notion of finite-dimensional wavelet systems on \mathbb{T}^2 , which employ a dilation operation induced by the Quincunx matrix. A theory of multiresolution analysis (MRA) is presented which includes the characterization and construction of MRA scaling functions in terms of low-pass filters. Orthonormal wavelet systems are constructed for any given MRA. Two general examples, based upon the classical Shannon and Haar wavelets, are presented and the approximation properties of the associated systems are studied.

4.1 Introduction

This work examines finite-dimensional systems of functions on the torus, \mathbb{T}^2 , which employ the basic tenets of wavelet theory: dilation and translation. The present study follows a similar analysis on the circle [2], \mathbb{T} , where dilation of $f \in L^2(\mathbb{T})$ was accomplished by a dyadic downsampling of the Fourier transform, i.e., $\widehat{Df}(k) = \widehat{f}(2k)$, $k \in \mathbb{Z}$. An obvious extension to $L^2(\mathbb{T}^2)$ would involve downsampling of the Fourier transform by $2I_2$; however, this choice fails to utilize the freedom provided by the move from one to two dimensions. Instead, the dilation operation considered here will be achieved through downsampling by a 2×2 matrix, A , satisfying

- A has integer entries.
- A has eigenvalues with modulus strictly greater than 1.
- A has determinant 2.

The first requirement is necessary for the downsampling operation $\widehat{Df}(k) = \widehat{f}(Ak)$, $k \in \mathbb{Z}^2$, to be well defined. The second condition ensures that repeated dilation of a function $f \in L^2(\mathbb{T}^2)$ will tend to a constant function in $L^2(\mathbb{T}^2)$. Finally, the third condition specifies that A should have minimal determinant. Indeed, if λ_1 and λ_2

B.D. Johnson (✉)
Saint Louis University, St. Louis, MO, USA
e-mail: brody@slu.edu

are the eigenvalues of A , then $|\det A| = |\lambda_1 \lambda_2|$ is an integer greater than 1. It is not difficult to see that if $\det A = 2$, then the trace of A will also be 2 under the above assumptions. A certain amount of the analysis will be independent of a specific choice for A . Nevertheless, A will hereafter denote the *Quincunx* dilation matrix,

$$A = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}.$$

This dilation is the composition of a rotation by $\frac{\pi}{4}$ with multiplication by $\sqrt{2}$ and, consequently, facilitates a natural geometric intuition. This discussion has focused on the role that A will play in the creation of a dilation operator. In the next section, the role played by A for translation will also be discussed.

4.2 Preliminaries

As mentioned in the previous section, the matrix A plays two roles in the proceeding theory, one dealing with dilation and another related to translation. In dilation, the Fourier transform of a function $f \in L^2(\mathbb{T}^2)$ will be downsampled over the subgroup $A\mathbb{Z}^2$ of \mathbb{Z}^2 . Translation will be considered over a discrete subgroup of \mathbb{T}^2 formed as a quotient of $A^{-j}\mathbb{Z}^2$ by \mathbb{Z}^2 , where $j > 0$ determines the *scale* or *resolution* of the translations being considered.

For a fixed integer $j > 0$, define the *lattice of order 2^j generated by A* , Γ_j , as the collection of 2^j distinct coset representatives of $A^{-j}\mathbb{Z}^2/\mathbb{Z}^2$. It will be assumed that each element of Γ_j belongs to the rectangle $[0, 1) \times [0, 1)$. In the next section, a notion of shift-invariant spaces will be introduced that consists of functions in $L^2(\mathbb{T}^2)$ which are invariant under translation by elements of Γ_j .

Recall that the dilation operation induced by A downsamples the Fourier transform of $f \in L^2(\mathbb{T}^2)$ by A . This operation will be best understood through the quotient groups $\mathbb{Z}^2/B^j\mathbb{Z}^2$, where $B = A^*$. Consequently, define the *dual lattice of order 2^j ($j > 0$) generated by A* , Γ_j^* , as the collection of 2^j distinct coset representatives of $\mathbb{Z}^2/B^j\mathbb{Z}^2$ determined by the intersection $B^jR \cap \mathbb{Z}^2$, where $R = (-\frac{1}{2}, \frac{1}{2}] \times (-\frac{1}{2}, \frac{1}{2}]$. Because B has integer entries it follows that $B^jR \subseteq B^{j+1}R$, so Γ_j^* is a natural subset of Γ_{j+1}^* (See Fig. 4.1).

The following lemma summarizes several elementary, but useful facts about the matrices A and B as well as the lattices Γ_j and Γ_j^* .

Lemma 4.1. *Let A, B as above and let $j \geq 2$ be an integer.*

1. $\Gamma_j = \{A^{-1}\alpha + \alpha' : \alpha \in \Gamma_{j-1}, \alpha' \in \Gamma_1\}$.
2. $\Gamma_j^* = \{B\beta + \beta' : \beta \in \Gamma_{j-1}^*, \beta' \in \Gamma_1^*\}$.
3. $AB^{-1} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$.
4. $A^k B^\ell = B^\ell A^k, k, \ell \in \mathbb{Z}$.




Fig. 4.1 The dual lattices Γ_3^* , Γ_4^* , and Γ_5^*

Another important feature of the lattices, Γ_j^* , is their behavior under multiplication by A . Hence, let $d_j : \Gamma_j^* \rightarrow \Gamma_j^*$ be the mapping defined by $\alpha \mapsto A\alpha$. For $1 \leq k \leq j-1$, let \mathcal{B}_k denote the kernel of d_j^k , i.e.,

$$\mathcal{B}_k = \{\beta \in \Gamma_j^* : d_j^k(\beta) = 0\}.$$

The following proposition provides two useful characterizations of \mathcal{B}_k .

Proposition 4.1. *Let $j \in \mathbb{N}$, $j \geq 2$ and $1 \leq k \leq j-1$.*

1. $\mathcal{B}_k = B^{j-k}\Gamma_j^*$, i.e., d_j^k is a 2^k -to-1 mapping.
2. $\mathcal{B}_k = \left\{ \sum_{\ell=1}^k b_\ell A^{j-\ell} \beta_1 : b_\ell \in \{0, 1\} \right\}$, where β_1 is the nonzero element of Γ_1^* .

Proof. To demonstrate the first claim, let $\beta \in \Gamma_j^*$ and assume that $d_j^k(\beta) = 0$. Hence, $A^k \beta \in B^j \mathbb{Z}^2$ or $A^k B^{-k} B^k \beta \in B^j \mathbb{Z}^2$. Now, since powers of A and B commute and AB^{-1} is a rotation, it follows that $B^k \beta \in B^j \mathbb{Z}^2$ and thus $\beta \in B^{j-k} \mathbb{Z}^2$. Likewise if $\beta \in B^{j-k} \mathbb{Z}^2$, then $A^k \beta \in B^j \mathbb{Z}^2$ and, thus, $d_j^k(\beta) = 0$.

To prove the second claim observe first that

$$A^k \sum_{\ell=1}^k b_\ell A^{j-\ell} \beta_1 = \sum_{\ell=1}^k b_\ell A^{k-\ell} A^j \beta_1 = 0,$$

because $A^{k-\ell}$ has integer entries and $A^j \beta_1 \in B^j \mathbb{Z}^2$. To see that the 2^k elements are unique, assume that

$$\sum_{\ell=1}^k b_\ell A^{j-\ell} \beta_1 = \sum_{\ell=1}^k b'_\ell A^{j-\ell} \beta_1,$$

or, equivalently, that

$$\sum_{\ell=1}^{k-1} (b_\ell - b'_\ell) A^{j-\ell} \beta_1 = (b'_k - b_k) A^{j-k} \beta_1.$$

Thus, the left- and right-hand quantities lie in the intersection of $A^{j-1}\Gamma_j^*$ and $A^{j-k}\Gamma_j^*$. However, $A^{j-k}\beta_1 \notin A^{j-1}\Gamma_j^*$, so it follows each expression must equal zero, i.e., $b_k = b'_k$. This argument may be repeated to show that $b_\ell = b'_\ell$, $1 \leq \ell \leq k$. \square

4.3 Shift-Invariant Spaces

This section introduces a notion of shift-invariant spaces for $L^2(\mathbb{T}^2)$ which make use of the lattices Γ_j , $j > 0$. The *translation operator generated by* $\alpha \in \Gamma_j$ will be denoted $T_\alpha : L^2(\mathbb{T}^2) \rightarrow L^2(\mathbb{T}^2)$ and is defined by

$$T_\alpha f(x) = f(x - \alpha), \quad x \in \mathbb{T}^2.$$

A shift-invariant space in this context will consist of a closed subspace V of $L^2(\mathbb{T}^2)$ with the property that $f \in V$ if and only if $T_\alpha f \in V$ for all $\alpha \in \Gamma_j$. Of course, if V is invariant under shifts in Γ_j , then V is also invariant under shifts in Γ_k , $1 \leq k \leq j$. This work shall focus attention on shift-invariant spaces generated by the Γ_j -translates of a single function.

Definition 4.1. Let $\phi \in L^2(\mathbb{T}^2)$. The *principal A -shift-invariant space of order 2^j generated by ϕ* , denoted $V_j(\phi)$, is the finite-dimensional subspace of $L^2(\mathbb{T}^2)$ spanned by the collection

$$X_j(\phi) = \{T_\alpha \phi : \alpha \in \Gamma_j\}. \quad (4.1)$$

A function in $V_j(\phi)$ is simply a linear combination of the Γ_j -translates of ϕ , which motivates the following definition. Let $\ell(\Gamma_j)$ denote the space of complex valued functions on Γ_j , with an analogous meaning for $\ell(\Gamma_j^*)$. Define $e_{j,\alpha} \in \ell(\Gamma_j^*)$, $j > 0$, $\alpha \in \Gamma_j$, by

$$e_{j,\alpha}(\beta) = \exp(2\pi i \langle \alpha, \beta \rangle), \quad \beta \in \Gamma_j^*.$$

Lemma 4.2. The collection $\{2^{-\frac{j}{2}} e_{j,\alpha}\}_{\alpha \in \Gamma_j}$ is an orthonormal basis for $\ell(\Gamma_j^*)$.

Proof. Given $\alpha', \alpha'' \in \Gamma_j$, the inner product of $e_{j,\alpha'}$ with $e_{j,\alpha''}$ can be expressed as

$$\langle e_{j,\alpha'}, e_{j,\alpha''} \rangle = \sum_{\beta \in \Gamma_j^*} \exp(2\pi i \langle \alpha, \beta \rangle),$$

where $\alpha = \alpha' - \alpha''$. If $\alpha = 0$ then the inner product is 2^j . However, if $\alpha \neq 0$, then there exists $\beta' \in \Gamma_j^*$ such that $\langle \alpha, \beta' \rangle \notin \mathbb{Z}$, in which case $\exp(2\pi i \langle \alpha, \beta' \rangle) \neq 1$. Since $\Gamma_j^* + \beta' \equiv \Gamma_j^*$, this leads to

$$\exp(2\pi i \langle \alpha, \beta' \rangle) \sum_{\beta \in \Gamma_j^*} \exp(2\pi i \langle \alpha, \beta \rangle) = \sum_{\beta \in \Gamma_j^*} \exp(2\pi i \langle \alpha, \beta \rangle),$$

which forces the sum, and hence the inner product, to be zero. \square

Recall that the Fourier transform of $f \in L^2(\mathbb{T}^2)$ is given by

$$\hat{f}(k) = \int_{\mathbb{T}^2} f(x) \exp(-2\pi i \langle x, k \rangle) dx, \quad k \in \mathbb{Z}^2.$$

Therefore, for $\alpha \in \Gamma_j$, $\widehat{T_\alpha f}(k) = \exp(-2\pi i \langle \alpha, k \rangle) \hat{f}(k)$, $k \in \mathbb{Z}^2$. The following definition adapts the familiar bracket product [1, 3] to the present context.

Definition 4.2. Let $f, g \in L^2(\mathbb{T}^2)$. The *A-bracket product of f and g of order 2^j* is the element of $\ell(\Gamma_j^*)$ defined by

$$[\hat{f}, \hat{g}]_{A^j}(\beta) = 2^j \sum_{k \in B^j \mathbb{Z}^2} \hat{f}(\beta + k) \overline{\hat{g}(\beta + k)}, \quad \beta \in \Gamma_j^*.$$

The bracket product so defined captures information about the inner products of f with the Γ_j -translates of g and can be effectively used to determine the frame properties of both principal and finitely-generated shift-invariant spaces. The following proposition, however, focuses on a characterization of orthonormal systems of Γ_j -translates.

Proposition 4.2. Let $f, g \in L^2(\mathbb{T}^2)$ and fix $\alpha \in \Gamma_j$. Then,

$$\langle T_\alpha f, g \rangle = 2^{-j} \langle [\hat{f}, \hat{g}]_{A^j}, e_{j, \alpha} \rangle.$$

In particular, $\langle T_\alpha f, g \rangle = \delta_{\alpha, 0}$, $\alpha \in \Gamma_j$, if and only if $[\hat{f}, \hat{g}]_{A^j}(\beta) = 1$, $\beta \in \Gamma_j^*$.

Proof.

$$\begin{aligned} \langle T_\alpha f, g \rangle &= \sum_{k \in \mathbb{Z}^2} \hat{f}(k) \overline{\hat{g}(k)} e^{-2\pi i \langle \alpha, k \rangle} \\ &= \sum_{\beta \in \Gamma_j^*} \sum_{k \in B^j \mathbb{Z}^2} \hat{f}(\beta + k) \overline{\hat{g}(\beta + k)} e^{-2\pi i \langle \alpha, \beta \rangle} \\ &= 2^{-j} \sum_{\beta \in \Gamma_j^*} [\hat{f}, \hat{g}]_{A^j}(\beta) \overline{e_{j, \alpha}(\beta)}. \end{aligned} \quad \square$$

The next step will be to incorporate dilation with the shift-invariant spaces examined here for the creation of multiresolution analyses.

4.4 A-Refinable Functions and Multiresolution Analysis

The goal of this section is to formulate a theory of multiresolution analysis on the torus making use of dilation by A and translations in Γ_j . The *dilation operator on $L^2(\mathbb{T}^2)$ induced by A* will be denoted $D : L^2(\mathbb{T}^2) \rightarrow L^2(\mathbb{T}^2)$ and is defined by

$$\widehat{Df}(k) = \hat{f}(Ak), \quad k \in \mathbb{Z}^2.$$

It follows that $DT_\alpha f = T_{B\alpha} Df$ for all $\alpha \in \Gamma_j$.

Definition 4.3. A function $\phi \in L^2(\mathbb{T}^2)$ is *A-refinable of order 2^j* if there exists a mask $c \in \ell(\Gamma_j)$ such that

$$D\phi = \sum_{\alpha \in \Gamma_j} c(\alpha) T_\alpha \phi. \quad (4.2)$$

If ϕ is refinable of order 2^j , then it follows that

$$\hat{\phi}(Ak) = \sum_{\alpha \in \Gamma_j} c(\alpha) \widehat{T_\alpha \phi}(k) = \sum_{\alpha \in \Gamma_j} c(\alpha) \overline{e_{j,\alpha}(k)} \hat{\phi}(k) = m(k) \hat{\phi}(k), \quad k \in \mathbb{Z}^2, \quad (4.3)$$

where $m = \sum_{\alpha \in \Gamma_j} c(\alpha) \overline{e_{j,\alpha}(\cdot)}$ is called the *filter* associated to ϕ . Note that each $k \in \mathbb{Z}^2$ belongs to the coset $\beta + B^j \mathbb{Z}^2$ of a unique element $\beta \in \Gamma_j^*$, i.e., $m \in \ell(\Gamma_j^*)$. The following lemma shows that the dilates of refinable functions are also refinable and provides a relationship between their filters.

Lemma 4.3. *If ϕ is refinable of order 2^j with filter $m \in \ell(\Gamma_j^*)$, then $D\phi$ is refinable of order 2^{j-1} with filter $m(A \cdot) \in \ell(\Gamma_{j-1}^*)$.*

Proof. Applying D to (4.2) and using the fact that $DT_\alpha = T_{B\alpha}D$ one finds that

$$D^2\phi = \sum_{\alpha \in \Gamma_j} c(\alpha) T_{B\alpha} D\phi.$$

This can be interpreted as a refinement equation for $D\phi$ of order 2^{j-1} , although the sum on the right includes duplicate representations of the elements of Γ_{j-1} . A straight-forward calculation shows that the above equation is equivalent to

$$D^2\phi = \sum_{\alpha \in \Gamma_{j-1}} \left(\sum_{\alpha' \in \Gamma_1} c(B^{-1}\alpha + \alpha') \right) T_\alpha D\phi.$$

In the Fourier domain this can be rewritten as $\widehat{D^2\phi} = \tilde{m} \widehat{D\phi}$, where $\tilde{m} \in \ell(\Gamma_{j-1}^*)$ is given by

$$\tilde{m}(\beta) = \sum_{\substack{\alpha \in \Gamma_{j-1} \\ \alpha' \in \Gamma_1}} c(B^{-1}\alpha + \alpha') \overline{e_{j-1,\alpha}(\beta)} = \sum_{\alpha \in \Gamma_j} c(\alpha) \overline{e_{j,\alpha}(A\beta)} = m(A\beta),$$

where $\beta \in \Gamma_{j-1}^*$. Recall that m is the filter associated to ϕ . □

As in [2], the usual notion of multiresolution analysis requires minor modifications for the torus.

Definition 4.4. A *multiresolution analysis (MRA) of order 2^j* ($j \in \mathbb{N}$) is a collection of closed subspaces of $L^2(\mathbb{T}^2)$, $\{V_k\}_{k=0}^j$, satisfying

1. For $1 \leq k \leq j$, $V_{k-1} \subseteq V_k$.
2. For $1 \leq k \leq j$, $f \in V_k$ if and only if $Df \in V_{k-1}$.

3. V_0 is the subspace of constant functions.
4. There exists a *scaling function* $\varphi \in V_j$ such that $X_k(2^{\frac{j-k}{2}} D^{j-k} \varphi)$ is an orthonormal basis for V_k , $0 \leq k \leq j$.

If φ is a scaling function for an MRA, then it is necessarily refinable and the filter associated to it by (4.3) is called a *low-pass filter* for φ (usually denoted by m_0). Moreover, the spaces V_k , $0 \leq k \leq j$, take the form $V_k(D^{j-k} \varphi)$ since $X_k(2^{\frac{j-k}{2}} D^{j-k} \varphi)$ is an orthonormal basis for V_k . The main results of this section follow. The first characterizes those refinable functions which are scaling functions for an MRA, while the second guarantees the existence of a scaling function given a suitable candidate filter, $m_0 \in \ell(\Gamma_j^*)$.

Theorem 4.1. *Suppose that $\varphi \in L^2(\mathbb{T}^2)$ is refinable of order 2^j ($j \in \mathbb{N}$) with $\hat{\varphi}(0) \neq 0$. Then φ is the scaling function of an MRA of order 2^j if and only if*

$$|m_0(\beta)|^2 + |m_0(\beta + B^{j-1}\beta_1)|^2 = 1, \quad \beta \in \Gamma_{j-1}^*, \quad (4.4)$$

and

$$[\hat{\varphi}, \hat{\varphi}]_{A^j}(\beta) = 1, \quad \beta \in \Gamma_j^*, \quad (4.5)$$

where β_1 is the nonzero element of Γ_1^* .

Proof. Assume that $\beta \in \Gamma_{j-1}^*$, then

$$\begin{aligned} [2^{\frac{1}{2}} \widehat{D\varphi}, 2^{\frac{1}{2}} \widehat{D\varphi}]_{A^{j-1}}(\beta) &= 2^j \sum_{k \in B^{j-1}\mathbb{Z}^2} |\widehat{D\varphi}(\beta + k)|^2 \\ &= 2^j \sum_{k \in B^{j-1}\mathbb{Z}^2} |\hat{\varphi}(\beta + k)|^2 |m_0(\beta + k)|^2 \\ &= 2^j \sum_{\substack{k \in B^j\mathbb{Z}^2 \\ \beta' \in \Gamma_1^*}} |\hat{\varphi}(\beta + B^{j-1}\beta' + k)|^2 |m_0(\beta + B^{j-1}\beta')|^2. \end{aligned}$$

Assume that (4.4) and (4.5) hold. It follows from the above calculation that

$$[2^{\frac{1}{2}} \widehat{D\varphi}, 2^{\frac{1}{2}} \widehat{D\varphi}]_{A^{j-1}}(\beta) = 1, \quad \beta \in \Gamma_{j-1}^*,$$

so $X_{j-1}(2^{\frac{1}{2}} D\varphi)$ is an orthonormal basis for its span. Moreover, Lemma 4.3 guarantees that the low-pass filter for $D\varphi$ satisfies (4.4), so, inductively, it follows that $X_k(2^{\frac{j-k}{2}} D^{j-k} \varphi)$ is an orthonormal basis for its span, $V_k(D^{j-k} \varphi)$, for each $0 \leq k \leq j$, guaranteeing MRA property 4. Properties 1 and 2 follow from the refinability of φ . Consider $2^{\frac{j}{2}} D^j \varphi$, which is refinable of order 2^0 with low-pass filter $m = 2^{\frac{j}{2}} m_0(A^j \cdot)$. Since Γ_0^* is the trivial quotient group, m is constant. In the Fourier domain, this means $\widehat{D^j \varphi}(k) = \hat{\varphi}(A^j k) = 2^{\frac{j}{2}} \hat{\varphi}(k)$. Recall that $\varphi \in L^2(\mathbb{T}^2)$, so this relation forces

$\widehat{D^j \varphi}(k) = 0$ unless $k = 0$, i.e., $V_0(\varphi)$ is the space of constant functions, justifying MRA property 3. (The fact that $\hat{\varphi}(0) \neq 0$ has also been used here.)

Conversely, assume that φ is the scaling function for an MRA. The orthonormality of $X_j(\varphi)$ is equivalent to (4.5). But $X_{j-1}(2^{\frac{1}{2}}D\varphi)$ is also an orthonormal collection and the calculation made at the beginning of the proof thus shows that

$$1 = [2^{\frac{1}{2}}\widehat{D\varphi}, 2^{\frac{1}{2}}\widehat{D\varphi}]_{A^{j-1}}(\beta) = |m_0(\beta)|^2 + |m_0(\beta + B^{j-1}\beta_1)|^2, \quad \beta \in \Gamma_{j-1}^*,$$

where β_1 is the nonzero element of Γ_1^* . Hence, (4.4) must hold, completing the proof. \square

The next theorem shows that the filter equation (4.4) is sufficient for the existence of a scaling function φ , provided that the candidate filter additionally satisfies $m(0) = 1$. A detailed discussion of examples will be postponed to Sect. 4.6, but it is fairly easy to come up with filters satisfying these requirements, e.g., define m_0 by

$$m_0(\beta) = \begin{cases} 1 & \beta \in \Gamma_{j-1}^* \\ 0 & \text{otherwise,} \end{cases} \quad \beta \in \Gamma_j^*.$$

The validity of this choice follows from Property 2 of Lemma 4.1, which implies $B^{j-1}\mathbb{Z}^2 = B^j\mathbb{Z}^2 \cup (B^{j-1}\beta_1 + B^j\mathbb{Z}^2)$. By the definition of Γ_{j-1}^* it follows that

$$\mathbb{Z}^2 = \bigcup_{\substack{\beta \in \Gamma_{j-1}^* \\ b \in \{0,1\}}} (\beta + bB^{j-1}\beta_1 + B^j\mathbb{Z}^2),$$

which shows that Γ_{j-1}^* and $\Gamma_{j-1}^* + B^{j-1}\beta_1$ form a partition of Γ_j^* .

Theorem 4.2. *Fix $j > 0$ and let $m_0 \in \ell(\Gamma_j^*)$ be a candidate low-pass filter satisfying (4.4) and $m_0(0) = 1$. Then, m_0 is the low-pass filter of a trigonometric polynomial scaling function of order 2^j .*

Proof. The proof will rest upon justification of a specific definition for an associated scaling function. The refinability will be accomplished by defining certain Fourier coefficients outside the lattice $A\mathbb{Z}^2$ (which by Property 3 of Lemma 4.1 is identical to $B\mathbb{Z}^2$) and extending using (4.2). Hence, let $\mathcal{B} = \Gamma_j^* \cap (B\mathbb{Z}^2)^c$ (\mathcal{B} should be regarded as a subset of \mathbb{Z}^2) and define $\varphi \in L^2(\mathbb{T}^2)$ as follows:

1. Let $\hat{\varphi}(0) = 2^{-\frac{j}{2}}$.
2. For $\beta \in \mathcal{B}$, let $\hat{\varphi}(\beta) = 2^{-\frac{j}{2}}$.
3. For $\beta \in \mathcal{B}$ and $1 \leq k \leq j-1$, let

$$\hat{\varphi}(A^k\beta) = \hat{\varphi}(\beta) \prod_{\ell=0}^{k-1} m_0(A^\ell\beta).$$

4. The remaining Fourier coefficients will be zero.

It is clear from the above definition that φ has finitely many nonzero Fourier coefficients, i.e., φ is a trigonometric polynomial. The refinability of φ with respect to the filter m_0 is inherent in the construction, provided that the *strands* defined in Step 3 terminate, i.e., there must exist k , with $1 \leq k \leq j-1$, such that $m_0(A^k \beta) = 0$. Proposition 4.1 implies that there are precisely two elements of Γ_j^* such that $A\beta = 0$, namely, 0 and $B^{j-1}\beta_1$. By definition, $A^j\beta = 0$ for all $\beta \in \Gamma_j^*$, so if $\beta \neq 0$ then the previous observation implies that $A^k\beta = B^{j-1}\beta_1$ for $1 \leq k \leq j-1$. Since (4.4) forces $m_0(B^{j-1}\beta_1) = 0$, this completes the proof of refinability. Therefore, in light of Theorem 4.1 it suffices to demonstrate (4.5), which will be accomplished through the following three steps.

1. It is a direct consequence of the definition above that (4.5) holds for each $\beta \in \mathcal{B} \cup \{0\}$. Hence, it remains only to demonstrate (4.5) for $\beta \in A\Gamma_j^* \setminus \{0\}$, a collection of $2^{j-1} - 1$ elements.
2. Proposition 4.1 explains that the mapping $\beta \mapsto A\beta$ is two-to-one, so the image of \mathcal{B} under multiplication by A^k has cardinality 2^{j-1-k} . Because $\mathcal{B} \subseteq \Gamma_j^* \setminus A\Gamma_j^*$, $A^{k_2}\mathcal{B}$ is disjoint from $A^{k_1}\mathcal{B}$ when $k_1 < k_2$. Considering $1 \leq k \leq j-1$ as in the construction above, the total number of unique elements of Γ_j^* belonging to $\{A^k\mathcal{B}\}_{k=1}^{j-1}$ is $2^{j-2} + 2^{j-3} + \dots + 2^1 + 1 = 2^{j-1} - 1$. None of these elements may belong to $\mathcal{B} \cup \{0\}$, so they are precisely the elements of $A\Gamma_j^* \setminus \{0\}$.
3. Let $\beta \in A\Gamma_j^* \setminus \{0\}$. Then $\beta = A^k(\beta' + \mathcal{B}_k)$ where $1 \leq k \leq j-1$ and $\beta' \in \mathcal{B}$. Moreover, using the fact that $\hat{\phi}(\beta') = 2^{-\frac{j}{2}}$ and Proposition 4.1,

$$\begin{aligned}
 [\hat{\phi}, \hat{\phi}]_{A^j}(\beta) &= \sum_{\gamma \in \mathcal{B}_k} \prod_{\ell=0}^k |m_0(A^\ell(\beta' + \gamma))|^2 \\
 &= |m_0(\beta')|^2 \sum_{\gamma \in \mathcal{B}_{k-1}} \prod_{\ell=0}^{k-1} |m_0(A^\ell(\beta' + \gamma))|^2 \\
 &\quad + |m_0(\beta' + B^{j-1}\beta_1)|^2 \sum_{\gamma \in \mathcal{B}_{k-1}} \prod_{\ell=0}^{k-1} |m_0(A^\ell(\beta' + \gamma + B^{j-1}\beta_1))|^2 \\
 &= \sum_{\gamma \in \mathcal{B}_{k-1}} \prod_{\ell=0}^{k-1} |m_0(A^\ell(\beta' + \gamma))|^2.
 \end{aligned}$$

This eventually reduces to the $k = 1$ case, which equals one by (4.4). \square

4.5 MRA Wavelets on the Torus

With the MRA theory of Sect. 4.4 it is a fairly straightforward task to devise a corresponding theory for MRA wavelets. An MRA of order 2^j consists of spaces $\{V_k\}_{k=0}^j$ with $V_{k-1} \subseteq V_k$, $1 \leq k \leq j$. In particular, V_j is a 2^j -dimensional subspace of

$L^2(\mathbb{T}^2)$ while V_0 is the one-dimensional subspace of constant functions. It is natural, therefore, to seek a wavelet system which provides an orthonormal basis for the orthogonal complement of V_0 in V_j , i.e., $V_j \ominus V_0$.

Definition 4.5. Let $\{V_k\}_{k=0}^j$ be an MRA of order 2^j . A function $\psi \in V_j$ is a *wavelet* for the MRA if the collection

$$\left\{ 2^{\frac{j-k}{2}} T_\alpha D^{j-(k+1)} \psi : 0 \leq k \leq j-1, \alpha \in \Gamma_k \right\}$$

is an orthonormal basis for $V_j \ominus V_0$.

The following theorem provides a construction of a wavelet for any MRA. The reader is reminded that β_1 denotes the nonzero element of Γ_1^* . Analogously, α_1 will denote the nonzero element of Γ_1 .

Theorem 4.3. Let φ be the scaling function of an MRA of order 2^j . Define ψ by

$$\hat{\psi}(k) = m_1(k) \hat{\varphi}(k), \quad k \in \mathbb{Z}^2,$$

where $m_1 \in \ell(\Gamma_j^*)$ is defined by

$$m_1(\beta) = \overline{m_0(\beta + B^{j-1}\beta_1)} \exp(2\pi i \langle A^{-(j-1)}\alpha_1, \beta \rangle). \quad (4.6)$$

Then, ψ is an orthonormal wavelet for the MRA.

Proof. The proof will establish an orthogonal decomposition of each space V_k , $1 \leq k \leq j$. By definition, $V_k = V_k(D^{j-k}\varphi)$, and the desired decomposition will have the form

$$V_k(D^{j-k}\varphi) = V_{k-1}(D^{j-k+1}\varphi) \oplus V_{k-1}(D^{j-k}\psi), \quad 1 \leq k \leq j.$$

In the wavelet literature, the spaces $V_{k-1}(D^{j-k}\psi)$ are often denoted W_k and one has the familiar expression $V_k = V_{k-1} \oplus W_k$, $1 \leq k \leq j$. The following calculation demonstrates the orthogonality of W_{j-1} and V_{j-1} . For each $\beta \in \Gamma_{j-1}^*$,

$$\begin{aligned} [2^{\frac{1}{2}} \widehat{D\varphi}, 2^{\frac{1}{2}} \hat{\psi}]_{A^{j-1}}(\beta) &= 2^j \sum_{k \in B^{j-1}\mathbb{Z}^2} \widehat{D\varphi}(\beta + k) \overline{\hat{\psi}(\beta + k)} \\ &= 2^j \sum_{k \in B^{j-1}\mathbb{Z}^2} |\hat{\varphi}(\beta + k)|^2 m_0(\beta + k) \overline{m_1(\beta + k)} \\ &= 2^j \sum_{\substack{k \in B^j\mathbb{Z}^2 \\ \beta' \in \Gamma_1^*}} |\hat{\varphi}(\beta B^{j-1}\beta' + k)|^2 m_0(\beta + B^{j-1}\beta') \overline{m_1(\beta + B^{j-1}\beta')} \\ &= m_0(\beta) \overline{m_1(\beta)} + m_0(\beta + B^{j-1}\beta_1) \overline{m_1(\beta + B^{j-1}\beta_1)} \\ &= 0, \end{aligned}$$

based upon (4.5), (4.6), and the fact that $\langle \alpha_1, \beta_1 \rangle = \frac{1}{2}$. The orthonormality of $X_{j-1}(2^{\frac{1}{2}}\psi)$ relies on a similar calculation, again for $\beta \in \Gamma_{j-1}^*$,

$$\begin{aligned}
 [2^{\frac{1}{2}}\hat{\psi}, 2^{\frac{1}{2}}\hat{\psi}]_{A^{j-1}}(\beta) &= 2^j \sum_{k \in B^{j-1}\mathbb{Z}^2} |\hat{\psi}(\beta + k)|^2 \\
 &= 2^j \sum_{k \in B^{j-1}\mathbb{Z}^2} |\hat{\phi}(\beta + k)|^2 |m_1(\beta + k)|^2 \\
 &= |m_1(\beta)|^2 + |m_1(\beta + B^{j-1}\beta_1)|^2 \\
 &= |m_0(\beta)|^2 + |m_0(\beta + B^{j-1}\beta_1)|^2 \\
 &= 1.
 \end{aligned}$$

The remainder of the proof stems from an induction argument. Lemma 4.3 implies that $2^{\frac{1}{2}}D\phi$ is refinable with filter $m_0(A\cdot)$, which satisfies the Γ_{j-1}^* equivalent of (4.4). Moreover, a calculation analogous to that in Lemma 4.3 shows that $2^{\frac{1}{2}}\widehat{D\psi} = m_1(A\cdot)2^{\frac{1}{2}}\widehat{D\phi}$, so the above calculations may be repeated at the next lower scale to prove that $X_k(2^{\frac{j-k}{2}}D^{j-(k+1)}\psi)$ is an orthonormal basis for W_k , $0 \leq k \leq j-1$. The orthogonality of W_{k_1} and W_{k_2} for $k_1 > k_2$ follows in the usual manner, i.e., $W_{k_2} \subseteq V_{k_1}$ which is orthogonal to W_{k_1} . \square

The last objective of this section is to examine the approximation provided by the wavelet systems considered in this work. The general approach mirrors that of [2]. If ψ is an orthonormal wavelet, then the system of functions in Definition 4.5 provides an orthonormal basis for $V_j \ominus V_0$ and together, with the constant function $D^j\phi$, can be used to approximate any $f \in L^2(\mathbb{T}^2)$. However, this is equivalent to considering the approximation of f by the collection $X_j(\phi)$. Consider the orthogonal projection $P_j : L^2(\mathbb{T}^2) \rightarrow V_j(\phi)$, given by

$$P_j f = \sum_{\alpha \in \Gamma_j} \langle f, T_\alpha \phi \rangle T_\alpha \phi.$$

In the Fourier domain this is equivalent to $\widehat{P_j f}(k) = [\hat{f}, \hat{\phi}]_{A^j}(k) \hat{\phi}(k)$, $k \in \mathbb{Z}^2$. For the purpose of this discussion, it suffices to consider f such that $\hat{f}(k) = \delta_{k,r}$, where $r, k \in \mathbb{Z}^2$. In this case,

$$[\hat{f}, \hat{\phi}]_{A^j}(\beta) = \begin{cases} 2^j \overline{\hat{\phi}(r)} & r \equiv \beta \pmod{B^j \mathbb{Z}^2} \\ 0 & r \not\equiv \beta \pmod{B^j \mathbb{Z}^2}, \end{cases} \quad \beta \in \Gamma_j^*,$$

so that

$$\widehat{P_j f}(k) = \begin{cases} 2^j \overline{\hat{\phi}(r)} \hat{\phi}(k) & r \equiv k \pmod{B^j \mathbb{Z}^2} \\ 0 & r \not\equiv k \pmod{B^j \mathbb{Z}^2}, \end{cases} \quad k \in \mathbb{Z}^2.$$

The squared error of approximation is thus given by

$$\begin{aligned}
 \|P_j f - f\|^2 &= \sum_{k \in \mathbb{Z}^2} |\widehat{P_j f}(k) - \hat{f}(k)|^2 \\
 &= (1 - 2^j |\hat{\phi}(r)|^2)^2 + 2^j |\hat{\phi}(r)|^2 2^j \sum_{\substack{k \in B^j \mathbb{Z}^2 \\ k \neq 0}} |\hat{\phi}(r+k)|^2 \\
 &= (1 - 2^j |\hat{\phi}(r)|^2)^2 + 2^j |\hat{\phi}(r)|^2 (1 - 2^j |\hat{\phi}(r)|^2) \\
 &= 1 - 2^j |\hat{\phi}(r)|^2,
 \end{aligned}$$

where fact that $[\hat{\phi}, \hat{\phi}]_{A^j} \equiv 1$ has been used to simplify the sum in the second term of the second line in this calculation. Define $E_j(k)$ by

$$E_j(k) = \sqrt{1 - 2^j |\hat{\phi}(k)|^2}, \quad k \in \mathbb{Z}^2.$$

Evidently, $E_j(k)$ represents the approximation error $\|P_j f - f\|$ when f is a trigonometric monomial with unit Fourier coefficient at $r \in \mathbb{Z}^2$. Observe that $E_j(k) = 0$ when $|\hat{\phi}(k)| = 2^{-\frac{j}{2}}$.

4.6 Examples

Since the systems considered in this work are finite-dimensional, proper examples should provide a well-defined MRA at any scale $j \geq 2$, hopefully leading to arbitrarily close approximation of functions in $L^2(\mathbb{T}^2)$. Moreover, given a low-pass filter satisfying $m_0(-\beta) = \overline{m_0(\beta)}$, $\beta \in \Gamma_j^*$ it is natural to expect a corresponding real-valued scaling function. The next proposition describes a modification of the construction in Theorem 4.2 that serves this purpose.

Proposition 4.3. *Let $m_0 \in \ell(\Gamma_j^*)$ be a low-pass filter satisfying (4.4) and such that $m_0(0) = 1$ and $m_0(-\beta) = \overline{m_0(\beta)}$, $\beta \in \Gamma_j^*$. Define ϕ as follows: (where $\beta \in \mathcal{B}$ should be regarded as an element of \mathbb{Z}^2)*

1. *Let $\hat{\phi}(0) = 2^{-\frac{j}{2}}$.*
2. *If $\beta, -\beta \in \mathcal{B}$, let $\hat{\phi}(\beta) = 2^{-\frac{j}{2}}$ and define*

$$\hat{\phi}(A^k \beta) = \hat{\phi}(\beta) \prod_{\ell=0}^{k-1} m_0(A^\ell \beta), \quad 1 \leq k \leq j-1.$$

3. *If $\beta \in \mathcal{B}$, but $-\beta \notin \mathcal{B}$, let $\hat{\phi}(\pm\beta) = 2^{-\frac{j+1}{2}}$ and define*

$$\hat{\phi}(\pm A^k \beta) = \hat{\phi}(\pm\beta) \prod_{\ell=0}^{k-1} m_0(\pm A^\ell \beta), \quad 1 \leq k \leq j-1.$$

4. *The remaining Fourier coefficients will be zero.*

Then φ is refinable with respect to m_0 and is a real-valued scaling function for an MRA of order 2^j .

Proof. The fact that φ is real-valued follows from the fact that the construction leads to the conjugate-symmetry $\hat{\varphi}(-k) = \overline{\hat{\varphi}(k)}$, $k \in \mathbb{Z}^2$. The m_0 -refinability is independent of the values of $\hat{\varphi}$ chosen at points in \mathcal{B} or $-\mathcal{B}$, provided that the refinement equation (4.2) is respected at the points in the A -orbit of such points. The fact that these orbits result in only finitely many nonzero Fourier coefficients for φ follows in exactly the same manner as in Theorem 4.2. Moreover, the argument given in the proof of Theorem 4.2 to demonstrate (4.5) requires only minor modification to account for the *splitting* of strands between $\pm\beta$ in Step 3, above. \square

Recall that the Shannon wavelet on \mathbb{R} is associated with an MRA consisting of band-limited subspaces of $L^2(\mathbb{R})$, with scaling function φ defined by $\hat{\varphi}(\xi) = \chi_{[-\frac{1}{2}, \frac{1}{2}]}(\xi)$ with corresponding low-pass filter $m_0(\xi) = \chi_{[-\frac{1}{4}, \frac{1}{4}]}$. An ideal analog of the Shannon MRA in this context should correspond to a low-pass filter which is symmetric about the origin and equal to the characteristic function of a set including 0. The following proposition describes a low-pass filter for each scale $j \geq 2$ which essentially captures these properties.

Proposition 4.4 (Shannon Filter). Fix $j \geq 2$ and let $S_j = \{\beta \in \Gamma_j^* : \beta, -\beta \in \Gamma_{j-1}^*\}$. The low-pass filter $m_0 \in \ell(\Gamma_j^*)$ defined by

$$m_0(\beta) = \begin{cases} 1 & \beta \in S_j \\ \frac{1}{\sqrt{2}} & \beta \in \Gamma_{j-1}^* \setminus S_j \quad \beta \in \Gamma_j^*, \\ \frac{1}{\sqrt{1 - |m_0(\beta - B^{j-1}\beta_1)|^2}} & \text{otherwise,} \end{cases}$$

satisfies (4.4) and is symmetric in the sense that $m_0(-\beta) = m_0(\beta)$, $\beta \in \Gamma_j^*$.

Proof. Recall that Γ_{j-1}^* and $\Gamma_{j-1}^* + B^{j-1}\beta_1$ form a partition of Γ_j^* , justifying the last part of the above definition. Hence, (4.4) is satisfied by construction.

The symmetry of m_0 requires attention to various cases. If $\beta \in S_j$, then $-\beta \in S_j$ and hence $m_0(\beta) = m_0(-\beta) = 1$. If $\beta \in \Gamma_{j-1}^* \setminus S_j$, then $m_0(\beta) = \frac{1}{\sqrt{2}}$. Moreover, $-\beta \notin \Gamma_{j-1}^*$ and, therefore, can be written as $-\beta = \beta' + B^{j-1}\beta_1$ for some $\beta' \in \Gamma_{j-1}^* \setminus S_j$. It follows that $m_0(\beta) = m_0(-\beta) = \frac{1}{\sqrt{2}}$. This demonstrates symmetry for all $\beta \in \Gamma_{j-1}^*$ and it now follows from (4.4) that m_0 is symmetric on all of Γ_j^* . \square

Figure 4.2 depicts the low-pass filter described by Proposition 4.4 for $j = 5$. Notice that the symmetry requirement, together with (4.4), makes it necessary to define $m_0(\beta) = m_0(-\beta) = \frac{1}{\sqrt{2}}$ for certain points in Γ_j^* . Figures 4.3 and 4.4 depict the scaling function and wavelet corresponding to the Shannon MRA.

The next proposition concerns the approximation of trigonometric polynomials provided by the Shannon MRA of order j .

Fig. 4.2 The low-pass filter m_0 of Proposition 4.4 with $j = 5$





Fig. 4.3 Graphs of the Shannon scaling function φ for $j = 5$: (a) $\varphi(x, y)$, $(x, y) \in \mathbb{T}^2$ and (b) surface plot of φ as a distortion of the torus




Fig. 4.4 Graphs of the wavelet function ψ corresponding to the Shannon MRA for $j = 5$: (a) $\psi(x, y)$, $(x, y) \in \mathbb{T}^2$ and (b) surface plot of ψ as a distortion of the torus

Proposition 4.5. *Let φ be the scaling function corresponding to the low-pass filter of Proposition 4.4 given by Proposition 4.3. If $j \geq 6 + \log_2 r^2$, then $E_j(k) = 0$ for all $k \in \{k = (k_1, k_2) : \max\{|k_1|, |k_2|\} \leq r\}$.*

Proof. Suppose that $\pm\beta \in \Gamma_{j-1}^*$. Then, Proposition 4.4 guarantees that $m_0(\pm\beta) = 1$. Moreover, if $\pm\beta \in \mathcal{B} \cup \{0\}$ then Proposition 4.3 implies that $\hat{\varphi}(\beta) = 2^{-\frac{j}{2}}$. Alternatively, if $\pm\beta \notin \mathcal{B} \cup \{0\}$, then $\pm\beta = \pm A^k \beta'$ for some β' such that $\pm\beta' \in \Gamma_{j-1}^*$ with $1 \leq k \leq j-1$. In this latter case, $|\hat{\varphi}(\beta')| = 2^{-\frac{j}{2}}$, while $m_0(A^\ell \beta') = 1$ for $1 \leq \ell \leq k-1$. The upshot of these observations is that if $\pm\beta \in \Gamma_{j-1}^*$, then $|\hat{\varphi}(\beta)| = 2^{-\frac{j}{2}}$. Therefore, $E_j(k) = 0$ whenever $\pm k \in \Gamma_{j-1}^*$.

Recall that $\Gamma_{j-1}^* = B^{j-1}R \cap \mathbb{Z}^2$, where $R = (-\frac{1}{2}, \frac{1}{2}] \times (-\frac{1}{2}, \frac{1}{2}]$. Instead consider $B^{j-1}R' \cap \mathbb{Z}^2$, where $R' = [-\frac{1}{4}, \frac{1}{4}] \times [-\frac{1}{4}, \frac{1}{4}]$, so that the set in question has symmetry about the origin. Observe that $B^{j-1}R'$ is a square with side-length $2^{\frac{j-3}{2}}$ centered at the origin and oriented with its corners either on the x, y -axes or along the lines $y = \pm x$. The former situation is the more constraining, but contains all $k = (k_1, k_2) \in \mathbb{Z}^2$ such that $\max\{|k_1|, |k_2|\} \leq \frac{1}{2} 2^{\frac{j-3}{2}} 2^{-\frac{1}{2}} = 2^{\frac{j}{2}-3}$. The claimed lower bound on j follows from this last calculation. \square

Another example important in the classical theory of MRAs is the Haar MRA. The Haar wavelet on \mathbb{R} is the product of an MRA whose component spaces consist of functions in $L^2(\mathbb{R})$ which are piecewise constant on certain dyadic intervals. The Haar scaling function φ is given by $\varphi = \chi_{[0,1]}$ and the corresponding low-pass filter is given by

$$m_0(\xi) = \frac{1}{2} (1 + \exp(-2\pi i \xi)).$$

Therefore, a natural counterpart to the Haar MRA should be associated with a conjugate-symmetric low-pass filter corresponding to a refinement involving just two translates from Γ_j . Moreover, assuming the first translate is zero, the nonzero translate should be as close to zero as possible. The following proposition describes a low-pass filter which meets these requirements.

Proposition 4.6 (Haar filter). *Fix $j \geq 2$. Define $m_0 \in \ell(\Gamma_j^*)$ by*

$$m_0(\beta) = \frac{1}{2} \left(1 + \exp(-2\pi i \langle A^{-(j-1)} \alpha_1, \beta \rangle) \right),$$

where α_1 is the nonzero element of Γ_1 . Then m_0 satisfies (4.4) with $m_0(0) = 1$ and is conjugate-symmetric, i.e., $m_0(-\beta) = \overline{m_0(\beta)}$, $\beta \in \Gamma_j^$.*

Proof. It is routine to verify that m_0 so defined is $B^j \mathbb{Z}^2$ -periodic and conjugate-symmetric with $m_0(0) = 1$. Observe that the filter may also be expressed as

$$m_0(\beta) = \cos(\pi \langle A^{-(j-1)} \alpha_1, \beta \rangle) \exp(-\pi i \langle A^{-(j-1)} \alpha_1, \beta \rangle).$$

Hence, the filter identity (4.4) follows from the calculation,

$$\begin{aligned}
 & |m_0(\beta)|^2 + |m_0(\beta + B^{j-1}\beta)|^2 \\
 &= \cos^2(\pi \langle A^{-(j-1)}\alpha_1, \beta \rangle) + \cos^2(\pi \langle A^{-(j-1)}\alpha_1, \beta + B^{j-1}\beta_1 \rangle) \\
 &= \cos^2(\pi \langle A^{-(j-1)}\alpha_1, \beta \rangle) + \cos^2\left(\pi \langle A^{-(j-1)}\alpha_1, \beta \rangle + \frac{\pi}{2}\right) \\
 &= \cos^2(\pi \langle A^{-(j-1)}\alpha_1, \beta \rangle) + \sin^2(\pi \langle A^{-(j-1)}\alpha_1, \beta \rangle) \\
 &= 1.
 \end{aligned}
 \quad \square$$

The final result of the section provides a somewhat coarse approximation result for the Haar MRA.

Proposition 4.7. *Let ϕ be the scaling function corresponding to the low-pass filter of Proposition 4.6 given by Proposition 4.3. Then for any $r \in \mathbb{Z}^2$,*

$$\lim_{j \rightarrow \infty} E_j(r) = 0.$$

Proof. Fix $r \in \mathbb{Z}^2$ and let J be the smallest positive integer such that $r \in B^J R \cap \mathbb{Z}^2$ (R as in the definition of Γ_j^*), so that $r = A^k \beta'$ for some $\beta' \in \mathcal{B}$ and $k \leq J$. For sufficiently large j , both β' and $-\beta'$ will belong to \mathcal{B} , so without loss of generality it may be assumed that $|\hat{\phi}(\beta)| = 2^{\frac{j}{2}}$. The construction of ϕ implies that

$$|\hat{\phi}(r)| = 2^{-\frac{j}{2}} \left| \prod_{\ell=0}^{k-1} m_0(A^\ell \beta') \right|$$

where

$$|m_0(A^\ell \beta')| = \cos(\pi \langle A^{-(j-1)}\alpha_1, A^\ell \beta' \rangle) = \cos(\pi \langle \alpha_1, (AB^{-1})^{j-1} A^{\ell+1-j} \beta' \rangle).$$

Notice that AB^{-1} is norm-preserving and $A^{\ell+1-j} \beta' \rightarrow 0$ as $j \rightarrow \infty$, which means that the terms in the above product tend to 1 as $j \rightarrow \infty$. Hence,

$$\lim_{j \rightarrow \infty} |\hat{\phi}(r)| = 2^{-\frac{j}{2}}$$

and $\lim_{j \rightarrow \infty} E_j(r) = 0$, concluding the proof. \square

Figure 4.5 depicts the modulus of the low-pass filter described by Proposition 4.6 for $j = 5$, while Figs. 4.6 and 4.7 depict the corresponding scaling function and wavelet for the Haar MRA.

Fig. 4.5 The modulus of the low-pass filter m_0 of Proposition 4.6 with $j = 5$





Fig. 4.6 Graphs of the Haar scaling function φ for $j = 5$: (a) $\varphi(x,y)$, $(x,y) \in \mathbb{T}^2$ and (b) surface plot of φ as a distortion of the torus

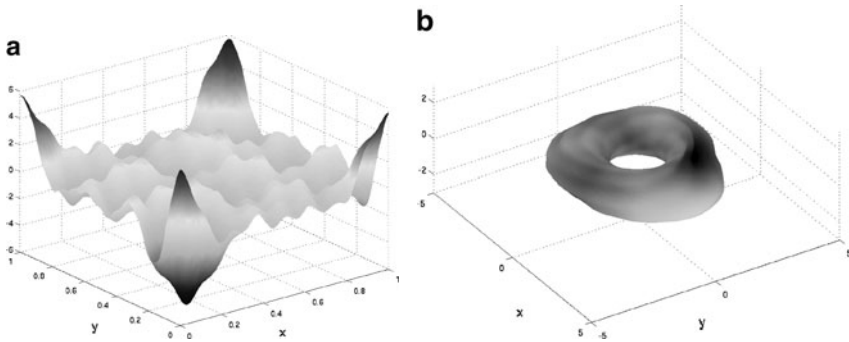


Fig. 4.7 Graphs of the wavelet function ψ corresponding to the Haar MRA for $j = 5$: (a) $\psi(x,y)$, $(x,y) \in \mathbb{T}^2$ and (b) surface plot of ψ as a distortion of the torus

References

1. C. de Boor, R. DeVore, and A. Ron, *The structure of finitely generated shift-invariant spaces in $L^2(\mathbb{R}^d)$* , J. Funct. Anal., **119**(1) (1995), 37–78.
2. B. D. Johnson, *A finite-dimensional approach to wavelet systems on the circle*, Glasnik Matematički, to appear.
3. A. Ron and Z. Shen, *Frames and stable bases for shift-invariant subspaces of $L^2(\mathbb{R}^d)$* , Canad. J. Math., **47** (1995), 1051–1094.

Chapter 5

Crystallographic Haar-Type Composite Dilation Wavelets

Jeffrey D. Blanchard and Kyle R. Steffen

Abstract An (a, B, Γ) composite dilation wavelet system is a collection of functions generating an orthonormal basis for $L^2(\mathbb{R}^n)$ under the actions of translations from a full rank lattice, Γ , dilations by elements of B , a subgroup of the invertible $n \times n$ matrices, and dilations by integer powers of an expanding matrix a . A Haar-type composite dilation wavelet system has generating functions which are linear combinations of characteristic functions. Krishtal, Robinson, Weiss, and Wilson introduced three examples of Haar-type (a, B, Γ) composite dilation wavelet systems for $L^2(\mathbb{R}^2)$ under the assumption that B is a finite group which fixes the lattice Γ . We establish that for any Haar-type (a, B, Γ) composite dilation wavelet, if B fixes Γ , known as the crystallographic condition, B is necessarily a finite group. Under the crystallographic condition, we establish sufficient conditions on (a, B, Γ) for the existence of a Haar-type (a, B, Γ) composite dilation wavelet. An example is constructed in \mathbb{R}^n and the theory is applied to the 17 crystallographic groups acting on \mathbb{R}^2 where 11 are shown to admit such Haar-type systems.

5.1 Introduction

Composite Dilation Wavelets (CDW) [14–16] are a recent generalization of wavelets [8, 18, among many] obtained from dilation by two countable sets of invertible matrices and translations by elements of a full rank lattice. As with classical wavelets, the theory of composite wavelets has seen early advances in the frequency domain and has now moved to the time domain with the introduction of Haar-type composite dilation wavelets [20]. An advantage of the composite dilation setting is the ability to introduce geometric (directional) information into the system as with the orientable oscillatory representation systems: ridglets [5], curvelets [6],

J.D. Blanchard (✉)

Department of Mathematics and Statistics, Grinnell College, Grinnell, Iowa, USA
e-mail: jeff@math.grinnell.edu

shearlets [12], etc. Naturally, it is important to know which composite dilation groups have a collection of functions with which one may construct a composite dilation wavelet system.

In this chapter, we consider composite dilation wavelets which arise from crystallographic groups and are linear combinations of characteristic functions, which we call crystallographic Haar-type composite dilation wavelets. We provide sufficient conditions under which such Haar-type composite dilation wavelet systems exist. Furthermore, under the crystallographic condition, we establish that the composite dilation group must be finite and provide an example of such a wavelet for an arbitrary dimension. The theory that is developed is applied to the crystallographic groups in \mathbb{R}^2 to show that 11 of these 17 groups admit Haar-type composite dilation wavelets.

5.1.1 Composite Dilation Wavelets

We denote the elements of \mathbb{R}^n , considered as column vectors, by letters from the Roman alphabet, usually x , k , or l . We denote elements of the Fourier domain by letters from the Greek alphabet and consider them to be row vectors, usually ξ , γ , or β . The letters f, ϕ, ψ shall be elements of the set of square integrable functions $L^2(\mathbb{R}^n)$, Φ, Ψ square integrable vector-valued functions, D, T are operators on functions, A, B, G subgroups of $GL_n(\mathbb{R})$ (the group of invertible matrices), and a, b, c matrices, often elements of a subgroup of $GL_n(\mathbb{R})$. Throughout, μ shall denote Lebesgue measure and the term *disjoint* is used in the Lebesgue measure sense. Occasionally, notation may differ but should be clear from context.

For an invertible matrix $c \in GL_n(\mathbb{R})$, define the *dilation of f by c* as $D_c f(x) = |\det c|^{-\frac{1}{2}} f(c^{-1}x)$. Let $\Gamma \subset \mathbb{R}^n$ be a full rank lattice, i.e. there exists $c \in GL_n(\mathbb{R})$ such that $\Gamma = c\mathbb{Z}^n$ and define the *translation of f by $k \in \Gamma$* as $T_k f(x) = f(x - k)$. With these two unitary operators, one constructs an affine system, $\mathcal{A}_C(\Psi) = \{D_c T_k \psi^\ell : c \in C, k \in \Gamma, 1 \leq \ell \leq L\}$, from a countable set of invertible matrices, $C \subset GL_n(\mathbb{R})$, a full rank lattice Γ , and a set of generating functions $\Psi = (\psi^1, \dots, \psi^L) \subset L^2(\mathbb{R}^n)$. Introduced by Guo, Labate, Lim, Weiss, and Wilson [14–16], affine systems with composite dilations are obtained when $C = AB$ is the product of two not necessarily commuting subsets of invertible matrices. In our setting, we further assume that $A = \{a^j : j \in \mathbb{Z}\}$ is a group generated by integer powers of an expanding matrix, a , and B is a subgroup of $GL_n(\mathbb{R})$, namely $\mathcal{A}_{aB\Gamma}(\Psi) = \{D_a^j D_b T_k \psi^\ell : j \in \mathbb{Z}, b \in B, k \in \Gamma, 1 \leq \ell \leq L\}$.

Definition 5.1 (composite dilation wavelet). $\Psi = (\psi^1, \dots, \psi^L) \subset L^2(\mathbb{R}^n)$ is a *composite dilation wavelet* if there exists an expanding matrix a , a group of invertible matrices B , and a full rank lattice Γ such that $\mathcal{A}_{aB\Gamma}(\Psi)$ is an orthonormal basis of $L^2(\mathbb{R}^n)$.

With a trivial composite dilation group, $B = \{I_n\}$ where I_n is the identity on \mathbb{R}^n , Definition 5.1 corresponds to the standard multiwavelet definition. Additionally, if $n = 1$, $\Gamma = \mathbb{Z}$, and $a = 2$, we recover the standard definition of a wavelet. For

classical wavelets, a performs expanding dilations of the wavelets while the translations act by “shifting” the supports of the wavelets. The introduction of the composite dilation group B in Definition 5.1 incorporates more geometric structure to the “shifts” of the wavelet supports. For CDW, a continues to perform the expanding dilations while dilations from B and translations from Γ act together as “shifts.”

In wavelet theory, an important property is the notion of a multiresolution analysis. Guo et al. showed that the MRA structure extends naturally to the setting of composite dilation wavelets [16].

Definition 5.2 (MRA). A nested sequence, $\{V_j\}_{j \in \mathbb{Z}}$, of closed subspaces of $L^2(\mathbb{R}^n)$ is an (a, B, Γ) -multiresolution analysis (MRA) if all of the following conditions are satisfied:

- (M1) $V_j \subset V_{j+1}$ where $V_j = D_a^{-j} V_0$.
- (M2) $\overline{\bigcup_{j \in \mathbb{Z}} V_j} = L^2(\mathbb{R}^n)$.
- (M3) $\bigcap_{j \in \mathbb{Z}} V_j = \{0\}$.
- (M4) There exists $\varphi \in V_0$ such that $\{D_b T_k \varphi : b \in B, k \in \Gamma\}$ is an orthonormal basis for V_0 .

A function φ which validates (M4) is called a *composite scaling function* in analogy to a scaling function associated with a classical MRA. Notice that the classical MRA is captured by this definition since a classical wavelet can be thought of as a composite dilation wavelet when B is the trivial group consisting only of the identity. In this case, we have that φ is a scaling function for the MRA if $\{T_k \varphi : k \in \Gamma\}$ is an orthonormal basis for V_0 .

Definition 5.1 is not the most general form of a composite dilation wavelet [14] but is appropriate for the current discussion. Some standard relaxations of the definition replace the requirement for an orthonormal basis with a Riesz basis or Parseval frame. Since the introduction of composite dilation wavelets, significant progress has been made in the theory of shearlets (e.g. [9, 12, 13, 21]) which are smooth functions generating a Parseval frame for $L^2(\mathbb{R}^2)$ and have compact support in frequency. The composite dilation group used to generate the shearlet systems is the group of integer shear matrices:

$$B = \left\{ \begin{pmatrix} 1 & j \\ 0 & 1 \end{pmatrix} : j \in \mathbb{Z} \right\}. \quad (5.1)$$

There is one example of an orthonormal CDW using the shear group (5.1) [16]. This is an example of a *minimally supported frequency* (MSF) composite dilation wavelet which consists of characteristic functions of compact sets in the frequency plane [16]. Many examples of MSF composite dilation wavelets were presented in the early composite dilation papers [14, 16], and it is known that every finite group B admits an MRA, MSF composite dilation wavelet [2, 3]. Here, we study which groups B admit an MRA CDW where the wavelets are characteristic functions in the time domain.

5.1.2 Haar-Type Composite Dilation Wavelets

In 1909, Alfred Haar [17] introduced an orthogonal system of functions for analyzing signals which is widely recognized as the first wavelet system. Haar's orthogonal function system neatly captures the information of a signal by storing averages and differences. The Haar system consists of two functions, one capturing averages and the other differences:

$$\varphi(x) = \chi_{[0,1]}(x), \quad (5.2)$$

$$\psi(x) = \chi_{[0,1/2]}(x) - \chi_{[1/2,1]}(x). \quad (5.3)$$

In the modern language of wavelets, φ is the Haar scaling function and ψ is the Haar wavelet of an MRA for $L^2(\mathbb{R})$. Clearly, these functions have excellent localization in the time domain. However, since these functions are not continuous, the frequency localization suffers.

The MSF composite dilation wavelets mentioned in Sect. 5.1.1 provide excellent frequency localization, but the desire to obtain significant time localization remains. Krishtal, Robinson, Weiss, and Wilson [20] present three examples of Haar-type composite dilation wavelets. A composite dilation wavelet is *Haar-type* if it carries two distinguishing features of the Haar wavelet, namely a multiresolution analysis structure and generating functions ψ^ℓ defined by linear combinations of characteristic functions.

A *Haar-type composite dilation wavelet* has an associated MRA and therefore an associated scaling function. Like the Haar wavelet, a Haar-type composite scaling function φ is a normalized characteristic function of a measurable set $R \subset \mathbb{R}^n$, $\varphi(x) = \frac{1}{\sqrt{\mu(R)}} \chi_R(x)$ where μ denotes Lebesgue measure. The Haar-type composite dilation wavelet generators are each a linear combination of a collection of measurable sets $R_{\ell,i} \subset \mathbb{R}^n$, i.e. there exist scalars $\alpha_{\ell,i} \in \mathbb{R}$ such that $\psi^\ell = \sum_i \alpha_{\ell,i} \chi_{R_{\ell,i}}$ for each $1 \leq \ell \leq L$. In this paper, we first study Haar-type composite dilation scaling functions and then obtain the wavelets from the unitary extension principle [4, 24].

5.1.3 The Crystallographic Condition

As mentioned above, the first examples of Haar-type composite wavelets were announced by Krishtal, Robinson, Weiss, and Wilson [20] under two assumptions: (a) B is a finite group and (b) the lattice Γ is invariant under the action of B , i.e. $B(\Gamma) = \Gamma$. In Sect. 5.2, we show that assumption (b) implies assumption (a), i.e. that B must be a finite group in a Haar-type composite dilation wavelet system. Assumption (b) is known as the crystallographic condition in group theory.

Definition 5.3 (Crystallographic Condition). A group of invertible matrices G and a full rank lattice Γ satisfy the *crystallographic condition* if Γ is invariant under the action of G , i.e. $G(\Gamma) = \Gamma$.

The crystallographic condition is usually applied only to subgroups of the orthogonal group acting on \mathbb{R}^n . For CDW, we apply the crystallographic condition more generally. When $B(\Gamma) = \Gamma$, the condition (M4) of Definition 5.2 can be viewed as requiring that the subspace V_0 is a $(B \ltimes \Gamma)$ -invariant space. For $(b_1, k_1), (b_2, k_2) \in B \ltimes \Gamma$, the group operation of the semi-direct product $B \ltimes \Gamma$ is given by

$$(b_1, k_1) \cdot (b_2, k_2) = (b_1 b_2, b_2^{-1} k_1 + k_2).$$

For more on this interpretation in CDW, see [16, 20].

However, the notion of a crystallographic group is used in the standard manner throughout our discussion.

Definition 5.4. Let G be a group of orthogonal transformations and Γ a full rank lattice for \mathbb{R}^n . G is a *crystallographic group* if Γ is a subgroup of G and the quotient group (called the point group) G/Γ is finite.

For example, the shear group (5.1) and \mathbb{Z}^2 satisfy the crystallographic condition (Definition 5.3) but the semi-direct product of the shear group and \mathbb{Z}^2 is not a crystallographic group since the quotient group is infinite. In Sect. 5.2.1, we prove that assuming the crystallographic condition for Haar-type CDW imposes that $B \ltimes \Gamma$ must be a crystallographic group, i.e. that B is finite. This implies, for instance, that no Haar-type CDW exist for the shear group (5.1).

The crystallographic groups are well-studied as they are the groups of symmetries of \mathbb{R}^n . The crystallographic condition derived its name from the important implications of crystallographic groups in the study of crystals. As a result, the crystallographic groups are especially well studied in \mathbb{R}^3 (for an introduction see [11]). Also, in \mathbb{R}^2 , the crystallographic groups are often referred to as wallpaper groups or plane symmetry groups and are commonly introduced in introductory algebra (for example, see [1, 25]). As indicated by the name, these wallpaper groups mathematically describe all repeated, congruent patterns that cover the plane and are often used for their geometric aesthetics.

Our goal is to construct *crystallographic Haar-type composite dilation wavelets* (CHCDW) which are Haar-type CDW where the semi-direct product of the composite dilation group B and the lattice Γ form a crystallographic group, i.e. $B \ltimes \Gamma$ satisfies Definition 5.4. To do so, we rely on two important measurable sets in \mathbb{R}^n , both of which are particular instances of a tiling set.

Definition 5.5 (Tiling Set). Let G be a group of invertible matrices and let R and W be measurable sets in \mathbb{R}^n . Then R is a *G -tiling set* of W if W is a disjoint union of the images of R under the action of G :

$$W = \bigcup_{g \in G} gR \quad \text{with} \quad \mu(g_1 R \cap g_2 R) = 0 \text{ for } g_1 \neq g_2 \in G.$$

When the set W is all of \mathbb{R}^n in Definition 5.5, we simply refer to R as a *G -tiling set* (with “of \mathbb{R}^n ” implied). While we utilize the notion of a tiling set in various capacities, there are two distinguished types of tiling sets for a crystallographic group.

A set $R \subset \mathbb{R}^n$ is called a *fundamental region* of the crystallographic group $B \ltimes \Gamma$ if R is a $(B \ltimes \Gamma)$ -tiling set. The second important subset of \mathbb{R}^n is a Γ -tiling set which tiles \mathbb{R}^n under the action of translations from the lattice Γ . In the construction of CHCDW, we will see how both of these sets play an important role. For example, the support set of a composite scaling function will always be a fundamental region for the group $B \ltimes \Gamma$ and therefore satisfy (M4) from Definition 5.2.

For composite dilation wavelets satisfying Definition 5.1, a crystallographic group must have a factorization into the semi-direct product $B \ltimes \Gamma$ in order to apply dilations from B and translations from Γ . In this case, since B is a subgroup of $B \ltimes \Gamma$, it is clear that $B(\Gamma) = \Gamma$ (see Lemma 5.2). This factorization into a semi-direct product is not necessary for producing wavelets in a more general setting. At the time of writing this article, the authors became aware of independent work by MacArthur and Taylor [23] in which a more general representation theoretic argument constructs Haar-type wavelet systems for $L^2(\mathbb{R}^2)$.

5.2 Crystallographic Composite Dilation Wavelets

Throughout this section, we study the implication of the crystallographic condition on Haar-type composite dilation wavelets. We first demonstrate that imposing the crystallographic condition, $B(\Gamma) = \Gamma$, on a Haar-type composite dilation system in fact forces us to use a group and lattice which form a crystallographic group, i.e. $B \ltimes \Gamma$ satisfies Definition 5.4. This is followed by establishing sufficient conditions for the existence of an (a, B, Γ) -MRA with a Haar-type composite scaling function. From this MRA and composite scaling function, we explore a method for constructing the CHCDW by first obtaining a low pass filter matrix and subsequently producing high pass filter matrices which define the composite dilation wavelets. For the remainder, we consider only Haar-type composite dilation wavelets. Therefore, $\varphi(x) = \frac{1}{\sqrt{\mu(R)}} \chi_R(x)$ for some measurable set $R \subset \mathbb{R}^n$.

5.2.1 Implications of the Crystallographic Condition

In this section, we investigate the impact of the crystallographic condition on CHCDW. We proceed through a sequence of lemmas demonstrating some geometric implications of imposing the crystallographic condition on our composite dilation system. The main result of this section, Theorem 5.1, states that the group B must be a finite subgroup of $\widetilde{SL}_n(\mathbb{R})$, the group of invertible matrices with determinant of absolute value 1, and therefore, $B \ltimes \Gamma$ is a crystallographic group.

Lemma 5.1. *If $\{D_b T_k \varphi(x) : b \in B, k \in \Gamma\}$ is an orthogonal system, then $\cup_{b \in B} bR$ is a disjoint union.*

Proof. This is an obvious consequence of orthogonality: Let $b_1, b_2 \in B$. If $b_1 \neq b_2$, then

$$0 = \mu(R) \langle D_{b_1} \varphi(x), D_{b_2} \varphi(x) \rangle = \int_{\mathbb{R}^n} \chi_{b_1 R}(x) \chi_{b_2 R}(x) d\mu = \mu(b_1 R \cap b_2 R).$$

□

A straightforward interpretation of Lemma 5.1 is that R must be a B -tiling set of $S = \cup_{b \in B} bR$ in order for $\varphi = \frac{1}{\sqrt{\mu(R)}} \chi_R$ to generate an orthonormal basis for the scaling space V_0 of the MRA in a Haar-type composite dilation wavelet system.

Lemma 5.2. *If $B(\Gamma) \subset \Gamma$ then $|\det(b)| = 1$ for all $b \in B$ and $B(\Gamma) = \Gamma$.*

Proof. Since $\Gamma = c\mathbb{Z}^n$ for some $c \in GL_n(\mathbb{R})$, there exists $0 \neq k \in \Gamma$ such that $\|k\|_2 \leq \|l\|_2$ for all $l \in \Gamma \setminus \{0\}$. Let $g \in GL_n(\mathbb{R})$ with $|\det(g)| < 1$. Then $0 < \|gk\|_2 < \|k\|_2$ and therefore $gk \notin \Gamma$ so $g \notin B$. Furthermore, $g^{-1} \notin B$ since B is a group. Finally, since I_n is an element of the group B , then $\Gamma \subset B(\Gamma)$ and, hence, $B(\Gamma) = \Gamma$. □

We now show that the disjoint union of the images of the scaling set R under the action of the group B is necessarily contained in a Γ -tiling set if B and Γ satisfy the crystallographic condition.

Lemma 5.3. *Suppose $\{D_b T_k \varphi(x) : b \in B, k \in \Gamma\}$ is an orthogonal system. If $B(\Gamma) = \Gamma$, then $\cup_{b \in B} bR$ is contained in a Γ -tiling set of \mathbb{R}^n .*

Proof. It suffices to show that for any $0 \neq k \in \Gamma$,

$$\mu((\cup_{b \in B} bR) \cap (\cup_{b \in B} bR) + k) = 0. \quad (5.4)$$

Let $b_1, b_2 \in B$. Then

$$\begin{aligned} \mu(b_1 R \cap (b_2 R + k)) &= \int_{\mathbb{R}^n} \chi_{b_1 R}(x) \chi_{(b_2 R + k)}(x) d\mu \\ &= \int_{\mathbb{R}^n} \chi_R(b_1^{-1} x) \chi_R(b_2^{-1} x - b_2^{-1} k) d\mu \\ &= \mu(R) \int_{\mathbb{R}^n} D_{b_1} \varphi(x) D_{b_2} T_{b_2^{-1} k} \varphi(x) d\mu. \end{aligned} \quad (5.5)$$

Since $B(\Gamma) = \Gamma$, there exists $l \in \Gamma$ such that $l = b_2^{-1} k \neq 0$. Thus, invoking (5.5) and the assumption that $\{D_b T_k \varphi(x) : b \in B, k \in \Gamma\}$ is an orthogonal system, we have

$$\mu(b_1 R \cap (b_2 R + k)) = \mu(R) \langle D_{b_1} \varphi(x), D_{b_2} T_l \varphi(x) \rangle = 0. \quad (5.6)$$

Since $b_1, b_2 \in B$ and $k \in \Gamma$ were all arbitrary, (5.4) follows directly from (5.6). □

The preceding lemmas combine to establish that the crystallographic condition on B and Γ imposes that B must be finite if B and Γ are to generate a Haar-type composite dilation wavelet system.

Theorem 5.1. *Suppose $\{D_b T_k \varphi(x) : b \in B, k \in \Gamma\}$ is an orthogonal system. If $B(\Gamma) = \Gamma$, then B is a finite subgroup of $\widetilde{SL}_n(\mathbb{R})$ and $B \ltimes \Gamma$ is a crystallographic group.*

Proof. From Lemma 5.3, there exists a Γ -tiling set, W , such that $\cup_{b \in B} bR \subset W$. Furthermore, since $\Gamma = c\mathbb{Z}^n$ for $c \in GL_n(\mathbb{R})$, it is clear that

$$\mu(\cup_{b \in B} bR) \leq \mu(W) = |\det(c)| < \infty. \quad (5.7)$$

Now, Lemmas 5.1 and 5.2 imply

$$\mu(\cup_{b \in B} bR) = \sum_{b \in B} |\det(b)| \mu(R) = |B| \mu(R). \quad (5.8)$$

Since $\mu(R) \neq 0$, it follows that $|B| < \infty$. That B is a subgroup of $\widetilde{SL}_n(\mathbb{R})$ follows from Lemma 5.2. Since B is finite, $B \ltimes \Gamma$ clearly satisfies Definition 5.4. \square

The proof of Theorem 5.1 provides a geometric insight into how the crystallographic condition imposes that the group B must be finite. For example, although the shear group (5.1) fixes \mathbb{Z}^2 , this is an infinite group and it is impossible to use the shear group to generate a Haar-type CDW. In related work, Houska [19] has proven a stronger result, namely that every composite dilation Bessel system with generating functions which are characteristic functions must have a finite group B . Houska's theorem directly implies that B must be finite in a Haar-type composite dilation wavelet system. While we were aware of this result, the preceding argument provides a direct, geometric connection between the crystallographic condition and the size of the group B , and the proof is distinct from Houska's.

5.2.2 Sufficient Conditions

We now turn our attention to establishing sufficient conditions on a , B , and Γ for the existence of a Haar-type (a, B, Γ) -composite dilation wavelet system for $L^2(\mathbb{R}^n)$. In Sect. 5.2.1, we saw that when B and Γ satisfy the crystallographic condition, the union of the images of the scaling set R under the action of B must be contained in a Γ -tiling set. Our first lemma states that if this union is a Γ -tiling set, then the scaling function $\varphi = \frac{1}{\sqrt{\mu(R)}} \chi_R$ generates an orthonormal basis for a $(B \ltimes \Gamma)$ -invariant space.

Lemma 5.4. *Suppose $B(\Gamma) = \Gamma$, R is a B -tiling set of S , and $\varphi = \frac{1}{\sqrt{\mu(R)}} \chi_R$. If $S = \cup_{b \in B} bR$ is a Γ -tiling set of \mathbb{R}^n , then $\{D_b T_k \varphi : b \in B, k \in \Gamma\}$ is an orthonormal system.*

Proof. Let $b_1, b_2 \in B$ and $k_1, k_2 \in \Gamma$. Since $B(\Gamma) = \Gamma$, there exist $l_1, l_2 \in \Gamma$ such that $l_1 = b_1 k_1$ and $l_2 = b_2 k_2$. By assumption, R is a B -tiling set of S and $b_1 R, b_2 R \subset S$ with S a Γ -tiling set of \mathbb{R}^n . Thus, for any $k \in \Gamma$ and any $b \in B$,

$$\mu(b_1R + k \cap b_2R + k) = 0 \text{ for } b_1 \neq b_2 \text{ and} \quad (5.9)$$

$$\mu(bR + k_1 \cap bR + k_2) = 0 \text{ for } k_1 \neq k_2. \quad (5.10)$$

Hence

$$\begin{aligned} \mu(R) \langle D_{b_1} T_{k_1} \varphi, D_{b_2} T_{k_2} \varphi \rangle &= \int_{\mathbb{R}^n} \chi_R(b_1^{-1}x - k_1) \chi_R(b_2^{-1}x - k_2) \, d\mu \\ &= \int_{\mathbb{R}^n} \chi_{(b_1R + b_1k_1)}(x) \chi_{(b_2R + b_2k_2)}(x) \, d\mu \\ &= \int_{\mathbb{R}^n} \chi_{(b_1R + l_1)}(x) \chi_{(b_2R + l_2)}(x) \, d\mu \\ &= \mu(b_1R + l_1 \cap b_2R + l_2). \end{aligned} \quad (5.11)$$

Now, from Lemma 5.2, $\mu(bR + k) = |\det(b)|\mu(R) = \mu(R)$ for every $b \in B$ and $k \in \Gamma$. Therefore, combining (5.9)–(5.11) yields

$$\langle D_{b_1} T_{k_1} \varphi, D_{b_2} T_{k_2} \varphi \rangle = \delta_{b_1, b_2} \delta_{k_1, k_2}$$

where δ is the usual Kronecker delta function. \square

Theorem 5.2 (Sufficient Conditions for an MRA). *Suppose $B(\Gamma) = \Gamma$. Let $\varphi = \frac{1}{\sqrt{\mu(R)}} \chi_R$ and $V_0 = \overline{\text{span}} \{D_b T_k \varphi : b \in B, k \in \Gamma\}$. Let $a \in GL_n(\mathbb{R})$ be an expanding matrix with $|\det(a)| = L + 1$.*

- (i) *Suppose $S = \cup_{b \in B} bR$ is a disjoint union and a Γ -tiling set of \mathbb{R}^n .*
- (ii) *Suppose $a\Gamma \subset \Gamma$ and a normalizes B (i.e. $aB = Ba$).*
- (iii) *Suppose there exist $b_0, \dots, b_L \in B$ and $k_0, \dots, k_L \in \Gamma$ such that*

$$aR = \bigcup_{i=0}^L (b_iR + b_i k_i). \quad (5.12)$$

Then the sequence of subspaces $\{V_j := D_a^{-j} V_0\}_{j \in \mathbb{Z}}$ is an MRA for $L^2(\mathbb{R}^n)$ and φ is a composite scaling function for this MRA.

Proof. From Lemma 5.4, condition (i) implies (M4). Since a is expanding, $\lim_{j \rightarrow \infty} \mu(a^j R) = \infty$. Therefore, the only square integrable function in every closed subspace V_j is the zero function which establishes (M3). As the piecewise constant functions are dense in $L^2(\mathbb{R}^n)$ and V_j is an approximation space spanned by $\{\chi_{a^{-j}(bR + bk)} : j \in \mathbb{Z}, b \in B, k \in \Gamma\}$, the closure of the union of all such approximation spaces is clearly $L^2(\mathbb{R}^n)$, verifying (M2).

Finally, we verify (M1). Since $V_j = D_a^{-j} V_0$, it suffices to show that $D_b T_k \varphi \in V_1$ for any $b \in B, k \in \Gamma$ as in the standard MRA arguments. From (5.12), $R = \cup_{i=0}^L a^{-1}(b_i R + b_i k_i)$ which is a disjoint union as argued in the proof Lemma 5.4.

Thus,

$$\chi_R(x) = \chi_{(\cup_{i=0}^L a^{-1}(b_i R + b_i k_i))}(x) = \sum_{i=0}^L \chi_R(b_i^{-1} a x - k_i),$$

and therefore

$$\varphi(x) = |\det(a)|^{-\frac{1}{2}} \sum_{i=0}^L D_a^{-1} D_{b_i} T_{k_i} \varphi(x). \quad (5.13)$$

Now for any $b \in B, k \in \Gamma$, (5.13) implies

$$\begin{aligned} D_b T_k \varphi(x) &= \varphi(b^{-1}x - k) = |\det(a)|^{-\frac{1}{2}} \sum_{i=0}^L D_a^{-1} D_{b_i} T_{k_i} \varphi(b^{-1}x - k) \\ &= |\det(a)|^{-\frac{1}{2}} \sum_{i=0}^L D_b T_k D_a^{-1} D_{b_i} T_{k_i} \varphi(x). \end{aligned} \quad (5.14)$$

With $a\Gamma \subset \Gamma$ and $B(\Gamma) = \Gamma$, there exist $l_i = b_i^{-1} a k \in \Gamma$. Hence, $T_k D_a^{-1} D_{b_i} = D_a^{-1} D_{b_i} T_{l_i}$. As a normalizes B , there exists $\tilde{b} \in B$ such that $ab^{-1} = \tilde{b}^{-1}a$ and therefore $D_b D_a^{-1} = D_a^{-1} D_{\tilde{b}}$. With these commuting relationships on the operators we can write (5.14) as

$$D_b T_k \varphi(x) = |\det(a)|^{-\frac{1}{2}} \sum_{i=0}^L D_a^{-1} D_{\tilde{b} b_i} T_{(l_i + k_i)} \varphi(x) \in V_1, \quad (5.15)$$

which verifies (M1). \square

To summarize, to obtain a Haar-type composite scaling function for an (a, B, Γ) -MRA, with $B \ltimes \Gamma$ a crystallographic group, it suffices to find a fundamental region of $B \ltimes \Gamma$ and an expanding matrix $a \in GL_n(\mathbb{R})$ such that $a\Gamma \subset \Gamma$, $aB = Ba$, and the action of a on the fundamental region, R , satisfies condition (iii).

5.2.3 Low Pass Filter Design

In the development of wavelet theory, properties of filters have played an important role. For example, Daubechies' construction of arbitrarily smooth compactly supported wavelets [7] was accomplished in part by factoring the filter associated with the scaling function. Moreover, the most common implementation algorithms rely completely on the filters. The notion of filters extends to MRA composite dilation wavelets [4, 16] and is useful for constructing wavelets from the composite scaling function.

As outlined in [4, 20], one may view the basis $\{D_b T_k \varphi\}$ of the space V_0 satisfying (M4) of Definition 5.2 as a Γ -invariant system with $|B|$ generating functions. Define $|B| = m + 1$ and fix an ordering to the elements of the composite dilation group, $B = \{b_0 = I_n, b_1, \dots, b_m\}$. Define the $m + 1$ functions

$$\varphi_s(x) := D_{b_s} \varphi(x) = \varphi(b_s^{-1}x) \text{ for } 0 \leq s \leq m. \quad (5.16)$$

Assuming the crystallographic condition, we have $D_b T_k = T_{bk} D_b$ with $bk \in \Gamma$, and hence, $\{D_b T_k \varphi : b \in B, k \in \Gamma\} = \{T_k D_b \varphi : k \in \Gamma, b \in B\} = \{T_k \varphi_s : k \in \Gamma, 0 \leq s \leq m\}$. Therefore, we may interpret the composite dilation scaling function as a multiscaling function with the group structure of B defining each of the scaling functions from the single function φ . Finally, we define the vector valued scaling function $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^{m+1}$ by

$$\Phi(x) = \begin{bmatrix} \varphi_0(x) \\ \vdots \\ \varphi_m(x) \end{bmatrix}. \quad (5.17)$$

For notational brevity, condition (iii) of Theorem 5.2 allows elements of the group B to appear in the set $\{b_0, \dots, b_L\}$ multiple times. Obviously, we could have $m < L$, requiring that for some $0 \leq i < i' \leq L$, $b_i = b_{i'}$. For the general construction of the low pass filter matrix, we introduce slightly more cumbersome notation by stating the equivalent condition

- (iii') Let $I \subset \{0, \dots, m\}$ be an index set and for each $i \in I$, define the index set P_i . Suppose there exist $\{b_i : i \in I\} \subset B$ and $\{k_{i,p} : p \in P_i\} \subset \Gamma$ with $\sum_{i \in I} |P_i| = L + 1$ such that

$$aR = \bigcup_{i \in I} \bigcup_{p \in P_i} (b_i R + b_i k_{i,p}). \quad (5.18)$$

Replacing (iii) with (iii') in Theorem 5.2 yields the same conclusion as the two conditions are equivalent. To track the commuting relationships, we further define the following index notation. Let $\tilde{s} \in \{0, \dots, m\}$ be defined by the normalizing relationship $b_{\tilde{s}} = a^{-1} b_s a$ and $s(i) \in \{0, \dots, m\}$ be defined by the right orbits of the elements of B , namely $b_{s(i)} = b_s b_i$. Since a normalizes the group B , both \tilde{s} and $s(i)$ are well-defined.

With this notation and assumption (iii'), following the proof of Theorem 5.2 and letting $k = 0$ in (5.15),

$$\begin{aligned} D_{b_s} \varphi(x) &= |\det(a)|^{-\frac{1}{2}} \sum_{i \in I} \sum_{p \in P_i} D_a^{-1} D_{b_{\tilde{s}} b_i} T_{k_{i,p}} \varphi(x) \\ &= |\det(a)|^{-\frac{1}{2}} \sum_{i \in I} \sum_{p \in P_i} D_a^{-1} D_{b_{\tilde{s}(i)}} T_{k_{i,p}} \varphi(x). \end{aligned} \quad (5.19)$$

In the multiscaling function notation (5.16), the commuting relationship $D_b T_k = T_{bk} D_b$ transforms (5.19) into

$$\begin{aligned} \varphi_s(x) &= |\det(a)|^{-\frac{1}{2}} \sum_{i \in I} \sum_{p \in P_i} D_a^{-1} T_{(b_{\tilde{s}(i)} k_{i,p})} \varphi_{\tilde{s}(i)}(x) \\ &= \sum_{i \in I} \sum_{p \in P_i} T_{(b_{\tilde{s}(i)} k_{i,p})} \varphi_{\tilde{s}(i)}(ax). \end{aligned} \quad (5.20)$$

Taking the Fourier transform, $\hat{f}(\xi) = \int_{\mathbb{R}^n} f(x) e^{-2\pi i \xi x} d\mu$, of (5.20) results in the equation

$$\hat{\phi}_s(\xi) = |\det(a)|^{-1} \sum_{i \in I} \sum_{p \in P_i} e^{-2\pi i (\xi a^{-1})(b_{\tilde{s}(i)} k_{i,p})} \hat{\phi}_{\tilde{s}(i)}(\xi a^{-1}). \quad (5.21)$$

From (5.21) we obtain a general construction for the low pass filter matrix.

Corollary 5.1. *Suppose conditions (i), (ii), and (iii') for Theorem 5.2 are satisfied with $I, P, \tilde{s}, s(i)$ defined above and let $|B| = m + 1$. Define the matrix $M_0(\xi)$ of size $(m + 1) \times (m + 1)$ whose $s, \tilde{s}(i)$ entry is*

$$[M_0(\xi)]_{s, \tilde{s}(i)} := \begin{cases} \frac{1}{|\det(a)|} \sum_{p \in P_i} e^{-2\pi i (\xi a^{-1})(b_{\tilde{s}(i)} k_{i,p})} & \text{for } i \in I, \\ 0 & \text{for } i \notin I. \end{cases} \quad (5.22)$$

Then $\hat{\Phi}(\xi a) = M_0(\xi) \hat{\Phi}(\xi)$.

Proof. Let $\hat{\Phi}(\xi)$ be the element-wise Fourier transform of $\Phi(x)$ defined by (5.17). The s th element of $\hat{\Phi}(\xi a)$ is then defined by (5.21), which is clearly the vector product of the s th row of $M_0(\xi)$ and $\hat{\Phi}(\xi)$ by (5.22). \square

5.2.4 Haar-Type CDW from the Composite Scaling Function

As with the traditional wavelets, the principal task in constructing an MRA wavelet system is obtaining a scaling function. Theorem 5.2 provides sufficient conditions for the construction of Haar-type composite scaling functions. In this section, we discuss two, closely related, systematic methods for obtaining appropriate composite dilation wavelet generating functions.

In [4], the first author and Krishtal established conditions on matrix filters which can be used to construct crystallographic composite dilation wavelets. This condition is the generalization of the unitary extension principle. We say that $M_0(\xi)$ is a *low pass filter* for Φ if $\hat{\xi}(\xi) = M_0(\xi a^{-1}) \hat{\xi}(\xi a^{-1})$ and that $M_\ell(\xi)$ is a *high pass filter* for $\Psi^\ell = (\psi_0^\ell, \dots, \psi_m^\ell)^t$ (defined similarly as (5.16)) if $\hat{\Psi}^\ell(\xi) = M_\ell(\xi a^{-1}) \hat{\Psi}^\ell(\xi a^{-1})$. Define $\hat{\mathbb{Z}}^n$ as the integer lattice of row vectors in the frequency domain, and let $\Gamma^* = \hat{\mathbb{Z}}^n c^{-1}$ denote the dual lattice to $\Gamma = c\mathbb{Z}^n$ so that $\gamma k \in \mathbb{Z}$ for any $\gamma \in \Gamma^*, k \in \Gamma$. The result from [4] is restated here with an appropriate notational modification.

Theorem 5.3 ([4]). *Let $B = \{b_0 = I, b_1, \dots, b_m\}$ and $\{V_j\}_{j=-\infty}^\infty$ be an MRA for $L^2(\mathbb{R}^n)$ with the associated composite scaling function Φ and the low pass filter $M_0(\xi)$. Let also $\{\beta_0, \dots, \beta_L\}$ be a full set of coset representatives of $\Gamma^* a^{-1} / \Gamma^*$, $\Psi^\ell = (\psi_0^\ell, \dots, \psi_m^\ell)^t$ with $\psi_s^\ell \in V_1$ for $s = 0, \dots, m$, and $M_\ell, 1 \leq \ell \leq L$, be the associated high pass filter. Then $\Psi = (\psi_s^\ell) \subset L^2(\mathbb{R}^n), 0 \leq s \leq m, 1 \leq \ell \leq L$, is a composite dilation wavelet for the (a, B, Γ) -MRA $\{V_j\}_{j=-\infty}^\infty$ if and only if*

$$\sum_{i=0}^L M_\ell(\xi + \beta_i) M_{\ell'}^*(\xi + \beta_i) = \delta_{\ell, \ell'} I_{m+1} \text{ for } 0 \leq \ell, \ell' \leq L. \quad (5.23)$$

With this result, the wavelets may be obtained by constructing appropriate high pass filter matrices from the low pass filter matrix. When $|\det(a)| = 2$, the full set of coset representatives is $\{0, \beta\}$ for some appropriate $\beta \in \mathbb{R}^n$. In this case, it was shown in [4] that when $M_0(\xi)M_0^*(\xi) = 0$, $M_1(\xi) = M_0(\xi + \beta)$ is a suitable high pass filter matrix which defines a composite dilation wavelet. A low pass filter for a crystallographic Haar-type composite scaling function (5.22) always satisfies the requirement $M_0(\xi)M_0^*(\xi) = 0$.

Alternatively, to define the CHCDW in the time domain, we might simply construct $L = |\det(a)| - 1$ functions $\psi^\ell, \ell = 1, \dots, L$, which are orthogonal to φ yet share the same support as φ . Sharing the same support imposes on these functions the familiar requirement $\int_{\mathbb{R}^n} \psi^\ell(x) d\mu = 0$. Since we are after Haar-type functions, we only need to alter the coefficients for the characteristic functions which define the dilated scaling function to construct an orthogonal function. A general approach is to find an orthogonal matrix of size $(L+1) \times (L+1)$ whose first row consists entirely of the constant $|\det(a)|^{-1/2}$. This first row represents the coefficients of the characteristic functions which define the dilated scaling function and ensures that the remaining row sums will all be zero.

Suppose $Q = (q_{\ell,t})_{\ell,t=0}^L$, with $q_{0,t} = 1/\sqrt{L+1} = |\det(a)|^{-1/2}$, is an orthogonal matrix. From (5.12), the dilated scaling function $\varphi(a^{-1}x)$ is supported on $L+1$ disjoint sets which we relabel $R_t, t = 0, \dots, L$. Then

$$D_a \varphi(x) = |\det(a)|^{-1/2} \varphi(a^{-1}x) = \frac{1}{\sqrt{\mu(R)}} \sum_{t=0}^L q_{0,t} \chi_{R_t}(x). \quad (5.24)$$

Now we define the wavelets by inserting coefficients obtained from the ℓ th row of Q .

$$D_a \psi^\ell(x) = |\det(a)|^{-1/2} \psi^\ell(a^{-1}x) = \frac{1}{\sqrt{\mu(R)}} \sum_{t=0}^L q_{\ell,t} \chi_{R_t}(x). \quad (5.25)$$

Then, with $\psi^0 = \varphi$, we have

$$\begin{aligned} \langle \psi^\ell, \psi^{\ell'} \rangle &= \int_{\mathbb{R}^n} \psi^\ell(x) \overline{\psi^{\ell'}(x)} d\mu \\ &= \frac{|\det(a)|}{\mu(R)} \int_{\mathbb{R}^n} \left[\sum_t q_{\ell,t} q_{\ell',t} \chi_{R_t}(ax) + \sum_{t \neq t'} q_{\ell,t} q_{\ell',t'} \chi_{R_t}(x) \chi_{R_{t'}}(ax) \right] d\mu \\ &= \frac{|\det(a)|}{\mu(R)} \sum_t q_{\ell,t} q_{\ell',t} \int_{\mathbb{R}^n} \chi_{R_t}(ax) d\mu \\ &= \frac{|\det(a)|}{\mu(R)} \mu(a^{-1}R) \sum_t q_{\ell,t} q_{\ell',t} \\ &= \delta_{\ell, \ell'} \end{aligned} \quad (5.26)$$

with the validity of this sequence of equalities following from (5.24) and (5.25), that R_t and $R_{t'}$ are disjoint for $t \neq t'$, and that Q is an orthogonal matrix.

Since each of the functions $\varphi_s(x)$ have disjoint supports and the functions $\psi_s^\ell(x) = D_{b_s} \psi^\ell(x)$ share these supports for every $s = 0, \dots, m$, the collection of functions $\{\psi_s^\ell : 0 \leq \ell \leq L, 0 \leq s \leq m\}$ is an orthonormal system. By simply taking the Fourier transform of Ψ^ℓ and extracting the associated high pass filters $M_\ell(\xi)$, Theorem 5.3 is satisfied and we have Haar-type composite dilation wavelets.

It is straightforward to observe that the traditional Haar scaling function (5.2) and wavelet (5.3) can be obtained in this fashion. For the standard Haar-wavelet, the group $B = \{1\}$ is the trivial group and the lattice is \mathbb{Z} . The dilation (matrix) is $a = 2$ and $\varphi(x) = \chi_{[0,1]}(x)$. Therefore,

$$\frac{1}{\sqrt{2}} \varphi\left(\frac{1}{2}x\right) = \frac{1}{\sqrt{2}} \chi_{[0,1]}(x) + \frac{1}{\sqrt{2}} \chi_{[1,2]}(x).$$

The dilated wavelet is obtained from the orthogonal matrix $Q = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ in the sense that

$$\frac{1}{\sqrt{2}} \psi\left(\frac{1}{2}x\right) = \frac{1}{\sqrt{2}} \chi_{[0,1]}(x) - \frac{1}{\sqrt{2}} \chi_{[1,2]}(x).$$

Since defining Haar-type composite dilation wavelets from the composite scaling function and an appropriate orthogonal matrix is effective in general, we define the following orthogonal matrices from which we define CHCDW in Sect. 5.3. The superscript represents the family of dilation matrices a with a common determinant, namely we have $Q^{|\det(a)|}$. While any such orthogonal matrix will suffice, we list some examples for expanding matrices whose determinants match those of the examples in Sect. 5.3.

$$Q^2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad Q^3 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 1 & 1 \\ \frac{\sqrt{6}}{2} & 0 & -\frac{\sqrt{6}}{2} \\ \frac{\sqrt{2}}{2} & -\sqrt{2} & \frac{\sqrt{2}}{2} \end{pmatrix} \quad Q^4 = \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & -1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \end{pmatrix}$$

$$Q^9 = \frac{1}{3} \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -\frac{1}{2} & -\frac{1}{2} & -1 & \frac{1}{2} & \frac{1}{2} & -1 & -1 & 2 \\ \frac{3}{2} & -\frac{3}{2} & 0 & \frac{3}{2} & -\frac{3}{2} & 0 & 0 & 0 & 0 \\ \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} & -\frac{2\sqrt{2}}{2} & -\frac{2\sqrt{2}}{2} & -\frac{2\sqrt{2}}{2} \\ \frac{\sqrt{3}}{2} & \frac{\sqrt{3}}{2} & -\frac{2\sqrt{3}}{2} & \frac{\sqrt{3}}{2} & \frac{\sqrt{3}}{2} & -\frac{2\sqrt{3}}{2} & 0 & 0 & 0 \\ \frac{\sqrt{6}}{2} & \frac{\sqrt{6}}{2} & \frac{\sqrt{6}}{2} & -\frac{\sqrt{6}}{2} & -\frac{\sqrt{6}}{2} & -\frac{\sqrt{6}}{2} & 0 & 0 & 0 \\ 0 & -\frac{\sqrt{6}}{2} & \frac{\sqrt{6}}{2} & 0 & \frac{\sqrt{6}}{2} & -\frac{\sqrt{6}}{2} & -\frac{\sqrt{6}}{2} & \frac{\sqrt{6}}{2} & 0 \\ -\frac{2\sqrt{30}}{10} & 0 & \frac{2\sqrt{30}}{10} & \frac{2\sqrt{30}}{10} & 0 & -\frac{2\sqrt{30}}{10} & \frac{\sqrt{30}}{10} & -\frac{3\sqrt{30}}{10} & \frac{2\sqrt{30}}{10} \\ -\frac{4\sqrt{5}}{10} & \frac{5\sqrt{5}}{10} & -\frac{\sqrt{5}}{10} & \frac{4\sqrt{5}}{10} & -\frac{5\sqrt{5}}{10} & \frac{\sqrt{5}}{10} & -\frac{8\sqrt{5}}{10} & \frac{4\sqrt{5}}{10} & \frac{4\sqrt{5}}{10} \end{pmatrix}$$

5.2.5 Example: $B \ltimes \Gamma = p4m$

In [20], a low pass filter matrix was constructed explicitly for the crystallographic group $B \ltimes \Gamma$ by observation of the action of a on a triangular fundamental region, R . In this example, $a = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$ is the quincunx matrix (which is a counter-clockwise rotation by $\pi/4$ and expansion by $\sqrt{2}$), $B = D_4$ is the group of the full symmetries of the square, and $\Gamma = \mathbb{Z}^2$. In this case $B \ltimes \Gamma = D_4 \ltimes \mathbb{Z}^2$ is the crystallographic group denoted $p4m$ (the interpretation of this group nomenclature is discussed in Sect. 5.3 or [25]).

Krishtal et al. fix the ordering of B as

$$b_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, b_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, b_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, b_3 = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix},$$

and $b_i = -b_{i-4}$ for $i = 4, 5, 6, 7$. The scaling function is defined as $\varphi = \frac{1}{\sqrt{\mu(R)}} \chi_R$ with the scaling set, R , defined as the isosceles right triangle with vertices, $\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \frac{1}{2} \\ 0 \end{pmatrix}, \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}$. One may easily verify conditions (i) and (ii) for Theorem 5.2. See Fig. 5.1.

To verify condition (iii'), we observe that $aR = b_1R \cup (b_6R + b_6\begin{pmatrix} -1 \\ 0 \end{pmatrix})$. Thus, the index sets are $I = \{1, 6\}$, $P_1 = \{\begin{pmatrix} 0 \\ 0 \end{pmatrix}\}$, and $P_6 = \{\begin{pmatrix} -1 \\ 0 \end{pmatrix}\}$. Moreover, we observe that this implies


$$D_a \varphi_0(x) = \frac{1}{\sqrt{2}} \varphi_0(a^{-1}x) = \frac{1}{\sqrt{2}} \varphi_1(x) + \frac{1}{\sqrt{2}} T_{\begin{pmatrix} 0 \\ 1 \end{pmatrix}} \varphi_6(x). \quad (5.27)$$

For s even, b_s is a rotation (or the identity), thus conjugation by a returns b_s . For s odd, b_s is a reflection and conjugation by a returns another reflection in B . Therefore,

$$\tilde{s} := \begin{cases} s & \text{for } s \text{ even,} \\ s + 2 \mod 8 & \text{for } s \text{ odd.} \end{cases} \quad (5.28)$$

Fig. 5.1 $B \ltimes \Gamma = p4m$:

(a) The support sets $R_s = b_s R, s = 0, \dots, 7$ for the scaling functions φ_s ; (b) The support sets of the dilated composite scaling function $\varphi_0(a^{-1}x)$ and composite wavelet $\psi_0(a^{-1}x)$ defined by (5.27) and (5.30)



Furthermore, examining the right orbits of $b_s \in B$ for $0 \leq s \leq m$,

$$s(i) := \begin{cases} s+i \pmod{8} & \text{for } s \text{ even,} \\ s-i \pmod{8} & \text{for } s \text{ odd.} \end{cases} \quad (5.29)$$

With (5.28) and (5.29), we can construct $M_0(\xi)$ for this example. Let $e(\alpha) = e^{2\pi i \alpha}$ and note that the indices of the $M_0(\xi)$ range from 0 to m . From (5.27), we can read off the entries of the first row (i.e. the row corresponding to $s = 0$). Since $\det(a) = 2$, the first row has $\frac{1}{2}$ in the second column and $\frac{1}{2}e(-\xi_2)$ in the seventh column, with the remaining entries of the row equal to zero. To further elucidate the notation, we will determine the entries of the second and third row, namely $s = 1, 2$. When $s = 1$, $\tilde{s} = 3$ and therefore $\tilde{s}(1) = 2$ and $\tilde{s}(6) = 5$. Since the entries of $M_0(\xi)$ are zero when $i \notin I$, the only nonzero entries appear in the columns indexed by 2 and 5, i.e. the third and sixth columns. Since P_1 contains only the zero vector, $[M_0(\xi)]_{1,2} = \frac{1}{2}$. Since $b_5 k_{6,0} = b_5 \begin{pmatrix} -1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, $[M_0(\xi)]_{1,5} = \frac{1}{2}e(-\xi_2)$. Likewise, with $s = 2$, $\tilde{s} = 2$ and $\tilde{s}(1) = 3, \tilde{s}(6) = 0$. Now $b_{\tilde{s}(1)} k_{1,0} = 0$ and $b_{\tilde{s}(6)} k_{6,0} = b_0 \begin{pmatrix} -1 \\ 0 \end{pmatrix} = \begin{pmatrix} -1 \\ 0 \end{pmatrix}$, therefore we have that $[M_0(\xi)]_{2,3} = \frac{1}{2}$ and $[M_0(\xi)]_{2,0} = \frac{1}{2}e(\xi_1)$. Continuing in this fashion, we reproduce the low pass filter matrix obtained in [20], namely

$$M_0(\xi) = \frac{1}{2} \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & e(-\xi_2) & 0 \\ 0 & 0 & 1 & 0 & 0 & e(-\xi_2) & 0 & 0 \\ e(\xi_1) & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & e(\xi_1) \\ 0 & 0 & e(\xi_2) & 0 & 0 & 1 & 0 & 0 \\ 0 & e(\xi_2) & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & e(-\xi_1) & 0 & 0 & 1 \\ 1 & 0 & 0 & e(-\xi_1) & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Now, since $\det(a) = 2$, we can construct the wavelets by either method detailed in Sect. 5.2.4. Since $\{\beta_0 = (0,0), \beta_1 = (1/2, 1/2)\}$ is a full set of coset representatives of $\hat{\mathbb{Z}}^2 a^{-1} / \hat{\mathbb{Z}}^2$, the matrix $M_1(\xi) = M_0(\xi + \beta_1)$ defines a high pass filter matrix:

$$M_1(\xi) = \frac{1}{2} \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & -e(-\xi_2) & 0 \\ 0 & 0 & 1 & 0 & 0 & -e(-\xi_2) & 0 & 0 \\ -e(\xi_1) & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & -e(\xi_1) \\ 0 & 0 & -e(\xi_2) & 0 & 0 & 1 & 0 & 0 \\ 0 & -e(\xi_2) & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & -e(-\xi_1) & 0 & 0 & 1 \\ 1 & 0 & 0 & -e(-\xi_1) & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Taking the first row of the inverse Fourier transform of $\hat{\Psi}^1(\xi) = M_1(\xi)\hat{\Psi}^1(\xi)$ yields the composite dilation wavelet

$$D_a\psi_0(x) = \frac{1}{\sqrt{2}}\psi_0(a^{-1}x) = \frac{1}{\sqrt{2}}\varphi_1(x) - \frac{1}{\sqrt{2}}T_{\begin{pmatrix} 0 \\ 1 \end{pmatrix}}\varphi_6(x). \quad (5.30)$$

Obviously, the same wavelet is obtained by selecting the coefficients from the second row of Q^2 and applying them to (5.27).

5.3 Crystallographic Haar-type Composite Dilation Wavelets

From the previous section, we have a systematic way to construct crystallographic Haar-type composite dilation wavelets. For a crystallographic group $B \ltimes \Gamma$ (B is finite), we select a fundamental region, R_0 , as the possible support set of the scaling function. To ensure that the characteristic function of this set χ_{R_0} will indeed define a scaling function, we must find an expanding matrix a which satisfies the conditions from Theorem 5.2. We then need to find appropriate coefficients to apply to the scaling equation to define the wavelets as outlined in Sect. 5.2.4. When this is possible, we have constructed a CHCDW. In Sect. 5.3.1, we demonstrate the existence of such a system in \mathbb{R}^n by constructing an example with a single wavelet generator. Section 5.3.2 answers the question almost completely for \mathbb{R}^2 by providing a catalog of the groups B which admit such Haar-type CDW.

5.3.1 A Crystallographic Haar-type Composite Dilation Wavelets for $L^2(\mathbb{R}^n)$

We begin by constructing a singly generated CHCDW in \mathbb{R}^n . Let $\Gamma = \mathbb{Z}^n$ and let B be the group generated by the n reflections through the hyperplanes perpendicular to the standard axes in \mathbb{R}^n . By construction, $B(\Gamma) = \Gamma$ and thus $B \ltimes \Gamma$ is a crystallographic group in \mathbb{R}^n . Now we center the unit hypercube, S , at the origin and take the fundamental region R_0 to be the hypercube with Lebesgue measure 2^{-n} in the all positive orthant, $R_0 = \{x_i \in [0, 1/2] : i = 1, \dots, n\}$. This is a fundamental region of $B \ltimes \Gamma$ and $S = \cup_{b \in B} bR_0$. Define the scaling function $\varphi(x) = 2^{-n/2}\chi_{R_0}(x)$.

For the expanding matrix, we select the modified permutation matrix

$$a = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 \\ 2 & 0 & 0 & \cdots & 0 & 0 \end{pmatrix}.$$




Fig. 5.2 (a) The support sets $R_s, s = 0, \dots, 7$ for the scaling functions φ_s described in Sect. 5.3.1 for \mathbb{R}^3 . (b) The support sets of the dilated composite scaling function $\varphi(a^{-1}x)$ and composite wavelet $\psi(a^{-1}x)$ described in Sect. 5.3.1 for \mathbb{R}^3

In this case, a normalizes B , and if B is ordered so that b_n is the reflection through the hyperplane perpendicular to the x_n -axis, then, with the n th standard basis vector denoted $k_n = (0, \dots, 0, 1)^t$,

$$aR_0 = R_0 \cup (b_n R_0 + k_n).$$

By Theorem 5.2, φ is a scaling function for an MRA in $L^2(\mathbb{R}^n)$. Since $|\det(a)| = 2$, we can construct the wavelet from the high pass filter $M_1(\xi) = M_0(\xi + \frac{1}{2}k'_n)$. In this case, $\psi(x) = \varphi(ax) - D_{b_n} T_{k_n} \varphi(ax)$. Figure 5.2 depicts the support sets for this construction in \mathbb{R}^3 . The matrix filter constructions for this example are more thoroughly treated in [4]. In \mathbb{R}^2 , this example corresponds to the plane crystallographic group pmm in the following section.

5.3.2 Crystallographic Haar-type Composite Dilation Wavelets for $L^2(\mathbb{R}^2)$

Krishtal et al. [20] showed the existence of simple, nonseparable, Haar-type CDW by providing three important examples. As mentioned in Sect. 5.2, they did so assuming the crystallographic condition. The preceding discussion has focused on fleshing out the theory of CHCDW and in this section we treat all crystallographic groups which act on \mathbb{R}^2 . There are 17 crystallographic groups for \mathbb{R}^2 which are often referred to as the wallpaper groups, the plane symmetry groups, or the plane crystallographic groups. These groups are well studied and an excellent introduction was written by Schattschneider [25].

In Sect. 5.3.2.1, we identify the 11 plane crystallographic groups which satisfy the conditions of Theorem 5.2. In Sect. 5.3.2.2, we construct support sets for the 11 groups identified in Sect. 5.3.2.1.

5.3.2.1 Plane crystallographic groups

The plane crystallographic groups admit 5 classes of lattices: parallelogram, rectangular, rhombic, square, and hexagonal. These lattices are a key ingredient in the proof that there are only 17 distinct plane crystallographic groups (up to isomorphism). Such a proof can be found throughout the literature, for example [1]. There are four basic isometric transformations in a plane crystallographic group: reflection, rotation, translation, and glide reflection. A glide reflection is a translation followed by a reflection across the translation axis. These groups have multiple standard notations assigned to them and we adopt the nomenclature from [25]. The only modifications to the following excerpt from [25] are our notion of a Γ -tiling set, the name of the axis (x_1 -axis), and the use of radians for angle measure.

The interpretation of the full international symbol (read left to right) is as follows: (1) letter p or c denotes primitive or centered Γ -tiling set; (2) integer n denotes the highest order of rotation; (3) symbol denotes a symmetry axis normal to the x_1 -axis: m (mirror) indicates a reflection axis, g indicates no reflection, but a glide-reflection, 1 indicates no symmetry axis; (4) symbol denotes a symmetry axis at angle α to x_1 -axis, with α dependent on n , the highest order of rotation: $\alpha = \pi$ for $n = 1$ or 2 , $\alpha = \frac{\pi}{4}$ for $n = 4$, $\alpha = \frac{\pi}{3}$ for $n = 3$ or 6 ; the symbols $m, g, 1$ are interpreted as in (3). No symbols in the third and fourth position indicate that the group contains no reflections or glide-reflections.

The 17 plane crystallographic groups are listed in Table 5.1, grouped by lattice type. Table 5.1 also has a column which specifies if the group can be factored in a semi-direct product $B \ltimes \Gamma$ with B finite. The groups which contain a glide reflection but no other reflection do not have such a factorization and are easily identified by the g which appears in their name. As mentioned briefly before, to fit the composite dilation structure, we form the affine system, $\mathcal{A}_{aB\Gamma}(\Psi) = \{D_a^j D_b T_k \Psi^\ell : j \in \mathbb{Z}, b \in B, k \in \Gamma\}$, which clearly requires the action of a group $B \ltimes \Gamma$ in the form of $D_b T_k$ for $(b, k) \in B \ltimes \Gamma$. Therefore, the four crystallographic plane groups without such

Table 5.1 The 17 plane crystallographic groups

Group	Lattice	Factorization	Group	Lattice	Factorization
$p1$	parallelogram	yes	$p4$	square	yes
$p2$	parallelogram	yes	$p4m$	square	yes
pm	rectangular	yes	$p4g$	square	no
pmm	rectangular	yes	$p3$	hexagonal	yes
pg	rectangular	no	$p31m$	hexagonal	yes
pmg	rectangular	no	$p3m1$	hexagonal	yes
pgg	rectangular	no	$p6$	hexagonal	yes
cm	rhombic	yes	$p6m$	hexagonal	yes
cmm	rhombic	yes			

Table 5.2 The factorization of the crystallographic groups into $B \ltimes \Gamma$. The generators for B are rotations and reflections: $\rho(\alpha)$ is a counterclockwise rotation by α ; $\sigma(\alpha)$ is a reflection across the line passing through the origin at an angle α with the x_1 -axis. The lattice $\Gamma = c\mathbb{Z}^2$ where c is chosen for the appropriate class of lattice: c_1, c_2 are arbitrary, real constants

Group Factorization					
Group	$B = \langle \text{generators} \rangle$	$\Gamma = c\mathbb{Z}^2$	Group	$B = \langle \text{generators} \rangle$	$\Gamma = c\mathbb{Z}^2$
$p1$	$\langle I \rangle$	$c \in GL_n(\mathbb{R})$	$p4$	$\langle \rho(\pi/2) \rangle$	$c = \begin{pmatrix} c_1 & 0 \\ 0 & c_1 \end{pmatrix}$
$p2$	$\langle \rho(\pi) \rangle$	$c \in GL_n(\mathbb{R})$	$p4m$	$\langle \sigma(0), \sigma(\pi/4) \rangle$	$c = \begin{pmatrix} c_1 & 0 \\ 0 & c_1 \end{pmatrix}$
pm	$\langle \sigma(0) \rangle$	$c = \begin{pmatrix} c_1 & 0 \\ 0 & c_2 \end{pmatrix}$	$p3$	$\langle \rho(2\pi/3) \rangle$	$c = c_1 \begin{pmatrix} \sqrt{3} & 0 \\ 1 & 2 \end{pmatrix}$
pmm	$\langle \sigma(0), \sigma(\pi/2) \rangle$	$c = \begin{pmatrix} c_1 & 0 \\ 0 & c_2 \end{pmatrix}$	$p31m$	$\langle \rho(2\pi/3), \sigma(\pi/6) \rangle$	$c = c_1 \begin{pmatrix} \sqrt{3} & 0 \\ 1 & 2 \end{pmatrix}$
cm	$\langle \sigma(0) \rangle$	$c = c_1 \begin{pmatrix} 2 & 1 \\ 0 & 1 \end{pmatrix}$	$p3m1$	$\langle \rho(2\pi/3), \sigma(0) \rangle$	$c = c_1 \begin{pmatrix} \sqrt{3} & 0 \\ 1 & 2 \end{pmatrix}$
cmm	$\langle \sigma(0), \sigma(\pi/2) \rangle$	$c = c_1 \begin{pmatrix} 2 & 1 \\ 0 & 1 \end{pmatrix}$	$p6$	$\langle \rho(\pi/3) \rangle$	$c = c_1 \begin{pmatrix} \sqrt{3} & 0 \\ 1 & 2 \end{pmatrix}$
			$p6m$	$\langle \sigma(0), \sigma(\pi/6) \rangle$	$c = c_1 \begin{pmatrix} \sqrt{3} & 1 \\ 0 & 2 \end{pmatrix}$

a factorization, pg , pmg , pgg , $p4g$, are not applicable to the composite dilation setting. Using a representation theoretic approach, MacArthur and Taylor constructed wavelets with the groups which do not fit the composite dilation setting [22, 23].

While $p3$ and $p31m$ admit a factorization into $B \ltimes \Gamma$, the nature of the fundamental region for these groups does not seem to allow the introduction of a matrix a which will satisfy the conditions of Theorem 5.2. Although a proof eludes us that no simple fundamental region can be paired with an expanding matrix a to satisfy the sufficient conditions of Theorem 5.2, an extensive search for a support set for a single scaling function has failed. For example, with a diagonal expanding matrix a , the expansion of the obvious fundamental regions always requires a translation of only half of one of the sets obtained by applying an element of B to this fundamental region. By permitting two scaling functions, MacArthur [22] identified appropriate support sets which produce the MRA obtained from crystallographic group for which $p3$ or $p31m$ are subgroups with index 2, namely $p6$ and $p6m$. Also, it seems plausible that support sets with fractal structure could be obtained by following the construction of Gröchenig and Madych [10]. Since we seek simple support sets, we omit this analysis.

5.3.2.2 Plane crystallographic Haar wavelets

We now employ the theory from Sect. 5.2 to construct CHCDW with simple support sets for the remaining 11 plane crystallographic groups. We treat 11 of the 17 since four groups (pg , pmg , pgg , $p4g$) have no factorization into $B \ltimes \Gamma$, and two

groups ($p3$, $p3m1$) do not seem to admit simple support sets. The process in the following discussion is to construct the support set of the scaling function, R_0 , and then demonstrate that the conditions of Theorem 5.2 are satisfied when an appropriate dilation matrix a is chosen. By presenting the support sets in Tables 5.3–5.7, we demonstrate that $S = \cup_{b \in B} bR_0$ is a disjoint union and a Γ -tiling set for \mathbb{R}^2 and that aR_0 is a union of translations of the sets $R_s = b_s R_0$. Throughout this section, we fix the ordering of the group $B = \{b_0 = I, b_1, \dots, b_m\}$ so that the sets $R_s = b_s R_0$ are enumerated in a counterclockwise fashion about the origin to form the Γ -tiling set S .

Table 5.3 Construction of support sets for CHCDW. The origin is depicted by a *bold dot*. The lattice basis elements, namely the columns of the matrix c for $\Gamma = c\mathbb{Z}^n$, are depicted by *arrows*. The columns with header aR_0 depicts condition (iii) of Theorem 5.2. The set aR_0 is outlined in *bold*

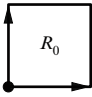
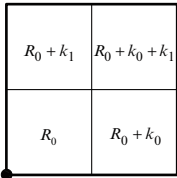
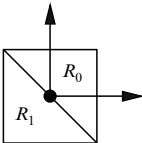
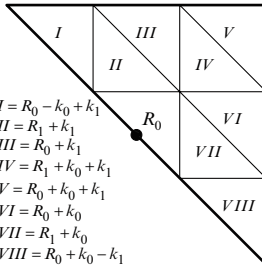
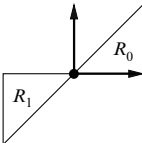
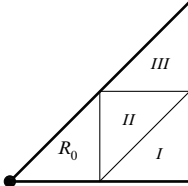
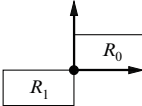
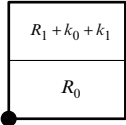
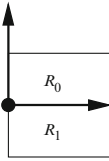
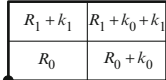
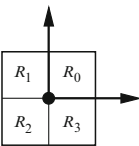
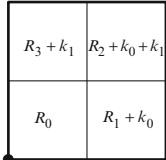
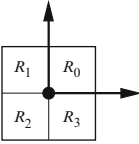
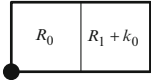
Square Lattice: $\Gamma = \mathbb{Z}^2$			
Group	a	Γ -tiling Set	aR_0
$p1$	$2I$		
$p2$	$3I$		 $I = R_0 - k_0 + k_1$ $II = R_1 + k_1$ $III = R_0 + k_1$ $IV = R_1 + k_0 + k_1$ $V = R_0 + k_0 + k_1$ $VI = R_0 + k_0$ $VII = R_1 + k_0$ $VIII = R_0 + k_0 - k_1$
$2I$			 $I = R_0 + k_0$ $II = R_1 + 2k_0 + k_1$ $III = R_0 + k_0 + k_1$
	$\begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}$		

Table 5.4 Construction of support sets for CHCDW. The origin is depicted by a *bold dot*. The lattice basis elements, namely the columns of the matrix c for $\Gamma = c\mathbb{Z}^n$, are depicted by *arrows*. The columns with header aR_0 depicts condition (iii) of Theorem 5.2. The set aR_0 is outlined in *bold*

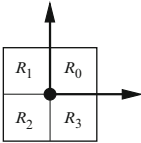
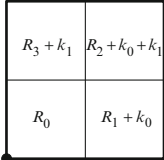
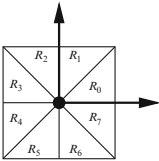
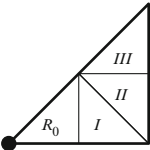
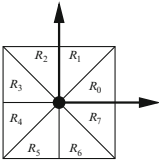
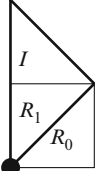
Square Lattice: $\Gamma = \mathbb{Z}^2$			
Group	a	Γ -tiling Set	aR_0
pm	$2I$		
pmm	$2I$		
	$\begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}$		

To form the $L = |\det(a)| - 1$ composite dilation wavelets, we follow the method outlined in Sect. 5.2.4. Select as coefficients for the characteristic functions the entries of the ℓ th row of the matrix $Q^{|\det(a)|}$ to form the dilated wavelet $D_a\psi^\ell(x)$. By applying dilations from B and taking the Fourier transform of the vector-valued wavelets $\Psi^\ell : \mathbb{R}^2 \rightarrow \mathbb{R}^{m+1}$, one may readily extract the high pass filter matrices $M_\ell(\xi)$ and verify the conditions of Theorem 5.3.

We catalog the CHCDW for these 11 plane crystallographic groups in Tables 5.3–5.7. The plane crystallographic groups are broken into lattice classes. The square lattice \mathbb{Z}^2 is by its nature also a parallelogram and rectangular lattice. These three lattice classes are all presented as a square lattice with the appropriate trivial generalizations (for diagonal dilation matrices) omitted. Tables 5.3–5.5 treat the six groups which admit a square lattice. The rhombic lattice applies to two groups, cm and cmm , which are presented in Table 5.6. Finally, $p3m1$, $p6$, and $p6m$ require translations from a hexagonal lattice and are treated in Table 5.7.

There are a few desirable, yet competing, properties to consider: the number of required composite dilation wavelet generating functions, separability of the scaling function, and connectedness of the interior of the Γ -tiling sets. Reducing the number of wavelet generators is certainly desirable and achieved by finding a matrix a with a small determinant. On the other hand, utilizing a dilation matrix of the form $a = \alpha I$ has advantages in implementation by eliminating a possible matrix multiplication. Clearly, this requires $\alpha^2 - 1$ composite dilation wavelet generating functions. Also, the additional structure of the composite dilation setting is most advantageous when the scaling function ϕ and the wavelets ψ^ℓ are not separable,

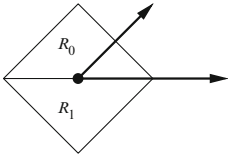
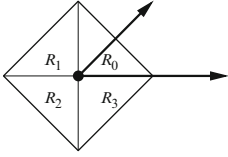
Table 5.5 Construction of support sets for CHCDW. The origin is depicted by a *bold dot*. The lattice basis elements, namely the columns of the matrix c for $\Gamma = c\mathbb{Z}^n$, are depicted by *arrows*. The columns with header aR_0 depicts condition (iii) of Theorem 5.2. The set aR_0 is outlined in *bold*; when R_0 is not contained in aR_0 , R_0 is included to depict the action of a on R_0

Square Lattice: $\Gamma = \mathbb{Z}^2$			
Group	a	Γ -tiling Set	aR_0
$p4$	$2I$		
$p4m$	$2I$		
	$\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$		
			$I = R_6 + k_1$

i.e. they are not products of univariate functions. It is often possible to reduce the number of generators by allowing separable scaling functions and wavelets. Finally, it may be advantageous to have a Γ -tiling set with connected interior. Again, by surrendering the requirement that the interior of $S = \cup_{b \in B} bR_0$ is connected, it is sometimes possible to reduce the determinant of a and therefore the necessary number of composite dilation wavelet generating functions.

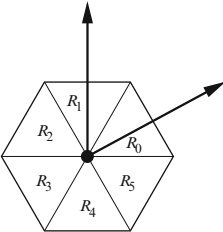
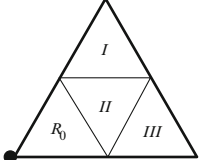
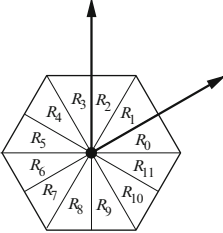
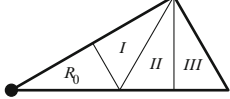
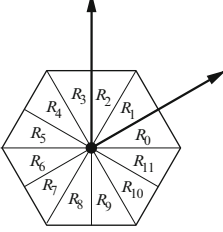
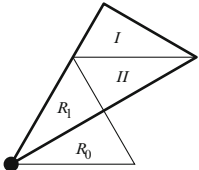
Tables 5.3–5.7 first treat a diagonal matrix $a = nI$ for $n \in \mathbb{Z}$. These are followed, in some cases, by examples where the properties from the previous paragraph are considered. When the support sets are rectangles, this is the case of a separable composite scaling function and composite dilation wavelet. The nonseparable scaling functions and composite dilation wavelets are supported on triangles. Four of

Table 5.6 Construction of support sets for CHCDW. The origin is depicted by a *bold dot*. The lattice basis elements, namely the columns of the matrix c for $\Gamma = c\mathbb{Z}^n$, are depicted by *arrows*. The columns with header aR_0 depicts condition (iii) of Theorem 5.2

Rhombic Lattice: $\Gamma = \begin{pmatrix} 2 & 1 \\ 0 & 1 \end{pmatrix} \mathbb{Z}^2$		
Group	a	Γ -tiling Set
cm	$3I$	 <div style="display: flex; justify-content: space-around; margin-top: 10px;"> <div style="text-align: left;"> $I = R_0 - k_0$ $II = R_1 - k_0 + k_1$ $III = R_0 - k_0 + k_1$ $IV = R_1 - k_0 + 2k_1$ </div> <div style="text-align: left;"> $V = R_0 - k_0 + 2k_1$ $VI = R_0 + k_1$ $VII = R_1 + k_1$ $VIII = R_0 + k_0$ </div> </div>
cmm	$3I$	 <div style="display: flex; justify-content: space-around; margin-top: 10px;"> <div style="text-align: left;"> $I = R_0 - k_0 + 2k_1$ $II = R_3 - k_0 + 2k_1$ $III = R_1 + k_1$ $IV = R_0 + k_1$ $V = R_2 + k_1$ $VI = R_3 + k_1$ $VII = R_1 + k_0$ $VIII = R_0 + k_0$ </div> <div style="text-align: left;"> $I = R_0 - k_0 + 2k_1$ $II = R_3 - k_0 + 2k_1$ $III = R_1 + k_1$ $IV = R_0 + k_1$ $V = R_2 + k_1$ $VI = R_3 + k_1$ $VII = R_1 + k_0$ $VIII = R_0 + k_0$ </div> </div>

these groups ($p1$, $p4m$, $p6$, $p6m$) have already been treated in the literature. Since B is the trivial group for $p1$, this group represents the case of standard wavelets in \mathbb{R}^2 . The case for a single Haar-type wavelet in \mathbb{R}^2 was characterized by Gröchenig and Madych [10] where the support set of the scaling function has a fractal boundary. The geometric complexity of the fractal boundaries does not meet our desire for simple support sets and this is omitted from the tables. Krishtal et al. presented three examples of simple Haar-type wavelets [20] for the groups $p4m$, $p6$, and $p6m$. The $p4m$ and $p6m$ examples from [20] appear as the second example for each in Tables 5.5 and 5.7, respectively. For $p6$, the expanding matrix used in [20] is the product of the counterclockwise rotation by $\pi/3$ and $2I$. Such combinations of a diagonal matrix and a generator for the composite dilation group B (see Table 5.2) are suitable expanding matrices for each of the eleven plane crystallographic groups in Tables 5.3–5.7.

Table 5.7 Construction of support sets for CHCDW. The origin is depicted by a *bold dot*. The lattice basis elements, namely the columns of the matrix c for $\Gamma = c\mathbb{Z}^n$, are depicted by *arrows*. The columns with header aR_0 depicts condition (iii) of Theorem 5.2. The set aR_0 is outlined in *bold*; when R_0 is not contained in aR_0 , R_0 is included to depict the action of a on R_0

Hexagonal Lattice: $\Gamma = \frac{3}{4} \begin{pmatrix} \sqrt{3} & 0 \\ 1 & 2 \end{pmatrix} \mathbb{Z}^2$			
Group	a	Γ -tiling Set	aR_0
$p6$ and $p3m1$	$2I$		 $I = R_2 + k_0$ $II = R_3 + k_0$ $III = R_4 + k_0$
$p6m$	$2I$		 $I = R_7 + k_0$ $II = R_8 + k_0$ $III = R_9 + k_0$
	$\frac{\sqrt{3}}{2} \begin{pmatrix} \sqrt{3} & -1 \\ 1 & \sqrt{3} \end{pmatrix}$		 $I = R_5 + k_0$ $II = R_6 + k_0$

Acknowledgements This work was partially supported by the University of Utah VIGRE program funded by NSF DMS grant number 0602219. JDB was partially funded as a VIGRE research assistant professor. KRS was funded through a year-long REU. The authors thank Guido Weiss, Ed Wilson, Ilya Krishtal, Keith Taylor, and Josh MacArthur for our rewarding conversations regarding CHCDW.

References

1. M. A. Armstrong. *Groups and symmetry*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1988.
2. J. D. Blanchard. Minimally supported frequency composite dilation wavelets. *J. Fourier Anal. App.* Online First, 2009, DOI:10.1007/s00041-009-9080-2.
3. J. D. Blanchard. Minimally supported frequency composite dilation Parseval frame wavelets. *J. Geo. Anal.*, 19(1):19–35, 2009.
4. J. D. Blanchard and I. A. Krishtal. Matricial filters and crystallographic composite dilation wavelets. submitted, 2009, http://www.math.grin.edu/~blanchaj/Research/MFCCDW_BIKr.pdf
5. E. Candès and D. L. Donoho. Ridgelets: a key to higher-dimensional intermittency? *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.*, 357(1760):2495–2509, 1999.
6. E. Candès and D. L. Donoho. Curvelets and curvilinear integrals. *J. Approx. Theory*, 113(1):59–90, 2001.
7. I. Daubechies. Orthonormal bases of compactly supported wavelets. *Comm. Pure Appl. Math.*, 41(7):909–996, 1988.
8. I. Daubechies. *Ten lectures on wavelets*, volume 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.
9. G. Easley, D. Labate, and W. Lim. Sparse directional image representations using the discrete shearlet transform. *Appl. Comput. Harmon. Anal.*, 25(1):25–46, 2008.
10. K. Gröchenig and W. R. Madych. Multiresolution analysis, Haar bases, and self-similar tilings of \mathbb{R}^n . *IEEE Trans. Inform. Theory*, 38(2, part 2):556–568, 1992.
11. L. C. Grove and C. T. Benson. *Finite reflection groups*, volume 99 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1985.
12. K. Guo, G. Kutyniok, and D. Labate. Sparse multidimensional representations using anisotropic dilation and shear operators. In *Wavelets and splines: Athens 2005*, *Mod. Methods Math.*, pages 189–201. Nashboro Press, Brentwood, TN, 2006.
13. K. Guo and D. Labate. Optimally sparse multidimensional representation using shearlets. *SIAM J. Math. Anal.*, 39(1):298–318, 2007.
14. K. Guo, D. Labate, W. Lim, G. Weiss, and E. N. Wilson. Wavelets with composite dilations. *Electron. Res. Announc. Amer. Math. Soc.*, 10:78–87, 2004.
15. K. Guo, D. Labate, W. Lim, G. Weiss, and E. N. Wilson. The theory of wavelets with composite dilations. In *Harmonic analysis and applications*, *Appl. Numer. Harmon. Anal.*, pages 231–250. Birkhäuser Boston, Boston, MA, 2006.
16. K. Guo, D. Labate, W. Lim, G. Weiss, and E. N. Wilson. Wavelets with composite dilations and their MRA properties. *Appl. Comput. Harmon. Anal.*, 20(2):202–236, 2006.
17. A. Haar. Zur theorie der orthogonalen Functionensysteme. *Math. Annalen*, 69:331–371, 1910.
18. E. Hernández and G. Weiss. *A first course on wavelets*. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL, 1996.
19. R. Houska. The nonexistence of shearlet-like scaling multifunctions that satisfy certain minimally desirable properties and characterizations of the reproducing properties of the integer lattice translations of a countable collection of square integrable functions. *Ph.D. dissertation*, 2009. Washington University in St. Louis.
20. I. Krishtal, B. Robinson, G. Weiss, and E. N. Wilson. Some simple Haar-type wavelets in higher dimensions. *J. Geom. Anal.*, 17(1):87–96, 2007.
21. G. Kutyniok and D. Labate. Resolution of the wavefront set using continuous shearlets. *Trans. Amer. Math. Soc.*, 361(5):2719–2754, 2009.
22. J. MacArthur. Compatible dilations and wavelets for the wallpaper groups. preprint, 2009.
23. J. MacArthur and K. Taylor. Wavelets from crystal symmetries. preprint, 2009.
24. A. Ron and Z. Shen. Affine systems in $L_2(\mathbb{R}^d)$: the analysis of the analysis operator. *J. Funct. Anal.*, 148(2):408–447, 1997.
25. D. Schattschneider. The plane symmetry groups: their recognition and notation. *Amer. Math. Monthly*, 85(6):439–450, 1978.

Chapter 6

From Full Rank Subdivision Schemes to Multichannel Wavelets: A Constructive Approach

Costanza Conti and Mariantonia Cotronei

Abstract In this paper, we describe some recent results obtained in the context of vector subdivision schemes which possess the so-called full rank property. Such kind of schemes, in particular those which have an interpolatory nature, are connected to matrix refinable functions generating orthogonal multiresolution analyses for the space of vector-valued signals. Corresponding multichannel (matrix) wavelets can be defined and their construction in terms of a very efficient scheme is given. Some examples illustrate the nature of these matrix scaling functions/wavelets.

6.1 Introduction

This paper deals with the construction of the proper “wavelet” tools for the analysis of functions which are *vector-valued*. Such type of functions arises naturally in many applications where the data to be processed are samples of vector-valued functions, even in several variables (e.g. electroencephalographic measurements, seismic waves, color images).

Up to now, most signal processing methods handle vector valued data component by component and thus ignore the possible relationship between several of the components. In reality, on the other hand, such measurements are usually related and thus should be processed as “complete” vectors. While it is obvious that for vector valued signals a standard scalar wavelet analysis does not take into account

C. Conti (✉)

Dipartimento di Energetica “Sergio Stecco”, Università di Firenze, Via Lombroso 6/17,
I-50134 Firenze, Italy
e-mail: costanza.conti@unifi.it

M. Cotronei (✉)

DIMET, Università “Mediterranea” di Reggio Calabria, Via Graziella,
I-89122 Reggio Calabria, Italy
e-mail: mariantonia.cotronei@unirc.it

the correlations among components, even an approach based on the well known *multiwavelets*, introduced in [7, 12] and widely studied (see [8] and the exhaustive references therein), is only apparently justified by the “vectorization” step required in order to apply them. Indeed, multiwavelets still generate multiresolution analyses for the space of scalar-valued functions.

In [13, 14] *matrix wavelets* have been proposed for the analysis of matrix valued signals while in [1] the concept of *multichannel wavelet* (MCW) and *multichannel multiresolution analysis* (MCMRA) have been introduced for processing vector-valued data. Actually, the two concepts look very similar: the crucial point of both is the existence of a matrix refinable function, which satisfies the so-called *full rank condition*. However, on one hand, the matrix analysis scheme proposed in [13, 14] coincides with the applications of a fixed number of MCW schemes, on the other hand, only a very special characterization of filters is proposed, which gives rise to interpolatory orthonormal matrix refinable functions. Thanks to this special structure, the corresponding wavelet filters are easily found as in the scalar situation. Nevertheless, in [13, 14] the authors do not provide either any constructive strategy nor any example. Some examples of 2×2 matrix wavelets are derived in [15], but they strongly rely on the sufficient condition for convergence described in [13], and belong again to a very special class. In passing, we also mention reference [6] since, in spite of the title, it actually contains examples of wavelet filters specifically designed for vector-fields.

In some recent papers [2–4] we investigated *full rank interpolatory schemes* and showed their connection to matrix refinable functions and multichannel wavelets. In particular, we proved that there is a connection between *cardinal* and *orthogonal* matrix refinable functions which is based on the *spectral factorization* of the (positive definite) symbol related to the cardinal function, thus, as a by-product, giving a concrete way to obtain orthogonal matrix scaling functions. The corresponding multichannel wavelets, whose existence is proved in [1], can be found in terms of a symbol which satisfies *matrix quadrature mirror filter* equations.

In this paper, we show how the solution of matrix quadrature mirror filter equations can be found. In particular, we give an efficient and constructive scheme which makes use again of spectral factorization techniques and of a matrix completion algorithm based on the resolution of generalized Bezout identities.

The paper is organized as follows. In Sect. 6.2, all the needed notation and definitions are set. Some introductory results on full rank vector subdivision schemes are also presented and their connection to orthogonal full rank refinable functions is shown. In Sect. 6.3, the concepts of MCMRA and MCW are revised while their orthogonal counterparts are considered next in Sect. 6.4. Section 6.5 deals with the construction of multichannel wavelets and an explicit algorithm for such construction is presented. In Sect. 6.6, starting from some interpolatory full rank positive definite symbols, the associated multichannel wavelets with two and three channels are constructed and shown. Some conclusions are derived in Sect. 6.7.

6.2 Notation and Basic Facts on Vector Subdivision Schemes

For $r, s \in \mathbb{N}$ we write a matrix $A \in \mathbb{R}^{r \times s}$ as $A = (a_{jk} : j = 1, \dots, r, k = 1, \dots, s)$ and denote by $\ell^{r \times s}(\mathbb{Z})$ the space of all $r \times s$ -matrix valued bi-infinite sequences, $\mathcal{A} = (A_j \in \mathbb{R}^{r \times s} : j \in \mathbb{Z})$. For notational simplicity we write $\ell^r(\mathbb{Z})$ for $\ell^{r \times 1}(\mathbb{Z})$ and $\ell(\mathbb{Z})$ for $\ell^1(\mathbb{Z})$ and denote vector sequences by lowercase letters.

By $\ell_2^{r \times s}(\mathbb{Z})$ we denote the space of all sequences which have finite 2 norm defined as

$$\|\mathcal{A}\|_2 := \left(\sum_{j \in \mathbb{Z}} \|A_j\|_2^2 \right)^{1/2}, \quad (6.1)$$

where $\|\cdot\|_2$ denotes the 2 operator norm for $r \times s$ matrices.

Moreover, $L_2^{r \times s}(\mathbb{R})$ will denote the Banach space of all $r \times s$ -matrix valued functions on \mathbb{R} with components in $L_2(\mathbb{R})$ and norm

$$\|F\|_2 = \left(\sum_{k=1}^s \sum_{j=1}^r \int_{\mathbb{R}} |F_{jk}(x)|^2 dx \right)^{1/2}. \quad (6.2)$$

For a matrix function and a matrix sequence of suitable sizes we introduce the *convolution* “ $*$ ” defined as,

$$(F * \mathcal{B}) := \sum_{k \in \mathbb{Z}} F(\cdot - k) B_k.$$

For two matrix functions $F, G \in L_2^{r \times s}(\mathbb{R})$, their inner product is defined as

$$\langle G, F \rangle = \int_{\mathbb{R}} G^H(x) F(x) dx.$$

and we have

$$\|F\|_2 = (\text{trace} \langle F, F \rangle)^{\frac{1}{2}}.$$

Let now \mathcal{A} be a *matrix sequence* $\mathcal{A} = (A_j \in \mathbb{R}^{r \times r} : j \in \mathbb{Z}) \in \ell^{r \times r}(\mathbb{Z})$. Let c be any bi-infinite *vector sequence* $c = (c_j \in \mathbb{R}^r : j \in \mathbb{Z})$. The *vector subdivision operator* $S_{\mathcal{A}}$, based on the matrix mask \mathcal{A} , acts on c by means of

$$S_{\mathcal{A}}c = \left(\sum_{k \in \mathbb{Z}} A_{j-2k} c_k : j \in \mathbb{Z} \right).$$

A *vector subdivision scheme* consists of iterative applications of a vector subdivision operator starting from an initial vector sequence $c \in \ell^r(\mathbb{Z})$, namely:

$$\begin{aligned} c^0 &:= c \\ c^n &:= S_{\mathcal{A}}c^{n-1} = S_{\mathcal{A}}^n c^0 \quad n \geq 1. \end{aligned}$$

It is L_2 -convergent if there exists a matrix-valued function $\mathbf{F} \in L_2^{r \times r}(\mathbb{R})$ such that

$$\lim_{n \rightarrow \infty} 2^{-n/2} \|\mu^n(\mathbf{F}) - S_{\mathcal{A}}^n \delta \mathbf{I}\|_2 = 0, \quad (6.3)$$

where the mean value operator at level $n \in \mathbb{N}$ is defined as

$$(\mu^n(\mathbf{F}))_j := 2^n \int_{2^{-n}(j+[0,1])} \mathbf{F}(t) dt, \quad j \in \mathbb{Z}.$$

The matrix-valued function \mathbf{F} associated in such a way with a convergent subdivision scheme is called the *basic limit function* and it is *refinable* with respect to \mathcal{A} , that is

$$\mathbf{F} = \mathbf{F} * \mathbf{A}(2 \cdot) = \sum_{j \in \mathbb{Z}} \mathbf{F}(2 \cdot - j) \mathbf{A}_j.$$

The *symbol* of a subdivision scheme $S_{\mathcal{A}}$ or of the mask \mathcal{A} is defined as:

$$\mathbf{A}(z) = \sum_{j \in \mathbb{Z}} \mathbf{A}_j z^j, \quad z \in \mathbb{C} \setminus \{0\},$$

and the *subsymbols*, $\mathbf{A}_{\varepsilon}(z) = \sum_{j \in \mathbb{Z}} \mathbf{A}_{\varepsilon+2j} z^j$, $z \in \mathbb{C} \setminus \{0\}$, $\varepsilon \in \{0, 1\}$, are related to the symbol by

$$\mathbf{A}(z) = \mathbf{A}_0(z^2) + z \mathbf{A}_1(z^2), \quad z \in \mathbb{C} \setminus \{0\}. \quad (6.4)$$

Next, we define the two $r \times r$ matrices

$$\mathbf{A}_{\varepsilon} := \mathbf{A}_{\varepsilon}(1) = \sum_{j \in \mathbb{Z}} \mathbf{A}_{\varepsilon+2j}, \quad \varepsilon \in \{0, 1\},$$

and their joint 1-eigenspace

$$\mathcal{E}_{\mathcal{A}} := \{\mathbf{y} \in \mathbb{R}^r : \mathbf{A}_0 \mathbf{y} = \mathbf{A}_1 \mathbf{y} = \mathbf{y}\}.$$

The *rank* of the mask \mathcal{A} or of the subdivision scheme $S_{\mathcal{A}}$ is the number

$$R(\mathcal{A}) := \dim \mathcal{E}_{\mathcal{A}}$$

satisfying $1 \leq R(\mathcal{A}) \leq r$ for convergent schemes. In particular, $S_{\mathcal{A}}$ is said to be of *full rank* if $R(\mathcal{A}) = r$.

The following proposition gives equivalent characterizations of full rank schemes [2]

Proposition 6.1. *Let $S_{\mathcal{A}}$ be a vector subdivision scheme with symbol $\mathbf{A}(z)$. Then the following statements are equivalent:*

- (i) $S_{\mathcal{A}}$ is of full rank.
- (ii) The symbol satisfies $\mathbf{A}(1) = 2\mathbf{I}$ and $\mathbf{A}(-1) = \mathbf{0}$, or, equivalently, there exists a matrix mask $\mathbf{B} \in \ell^{r \times r}(\mathbb{Z})$ such that $\mathbf{A}(z) = (z+1)\mathbf{B}(z)$, with $\mathbf{B}(1) = \mathbf{I}$ (scalar-like factorization).

- (iii) *The scheme preserves all vector constant data, i.e. whenever $\mathbf{c}_j = \mathbf{d}$, $j \in \mathbb{Z}$, for some $\mathbf{d} \in \mathbb{R}^r$, then also $(S_{AC})_j = \mathbf{d}$, $j \in \mathbb{Z}$.*
- (iv) *The basic limit function \mathbf{F} is a partition of the identity, i.e. $\sum_{j \in \mathbb{Z}} \mathbf{F}(\cdot - j) = \mathbf{I}$.*

A full rank vector subdivision scheme is *interpolatory* if it is characterized by the property

$$(S_{AC})_{2j} = \mathbf{c}_j, \quad j \in \mathbb{Z},$$

or, equivalently,

$$\mathbf{A}_{2j} = \delta_{j0} \mathbf{I}, \quad j \in \mathbb{Z}.$$

Its associated basic limit function, whenever it exists, turns out to be *cardinal*, i.e.,

$$\mathbf{F}(j) = \delta_{j0} \mathbf{I}, \quad j \in \mathbb{Z}.$$

The connection between interpolatory and full rank subdivision schemes is expressed by the following result [2]:

Proposition 6.2. *An interpolatory subdivision scheme with symbol $\mathbf{A}(z)$ is of full rank if and only if $\mathbf{A}(1) = 2\mathbf{I}$.*

6.3 Multichannel Multiresolution Analysis and Multichannel Wavelets

As in the scalar situation, a *multichannel multiresolution analysis* (MCMRA) in the space $L_2^r(\mathbb{R})$ of square integrable vector valued functions can be defined by a nested sequence $\cdots \subset V_{-1} \subset V_0 \subset V_1 \subset \cdots$ of closed subspaces of $L_2^r(\mathbb{R})$ with the properties that:

1. They are *shift invariant*

$$\mathbf{h} \in V_j \iff \mathbf{h}(\cdot + k) \in V_j, \quad k \in \mathbb{Z}.$$

2. They are *scaled* versions of each other

$$\mathbf{h} \in V_0 \iff \mathbf{h}(2^j \cdot) \in V_j.$$

$$3. \bigcap_{j \in \mathbb{Z}} V_j = \{(0, 0 \dots 0)^T\}, \quad \overline{\bigcup_{j \in \mathbb{Z}} V_j} = L_2^r(\mathbb{R}).$$

4. The space V_0 is generated by the integer translates of r *function vectors*, that is there exists $\mathbf{f}^i = (f_1^i, \dots, f_r^i)^T \in L_2^r(\mathbb{R})$, $i = 1, \dots, r$, such that

$$(i) \quad V_0 = \text{span} \{ \mathbf{f}^i(\cdot - k) : k \in \mathbb{Z}, i = 1, \dots, r \},$$

(ii) There exist two constants $K_1 \leq K_2$ such that

$$K_1 \left(\sum_{i=1}^r \sum_{k \in \mathbb{Z}} |c_k^i|^2 \right)^{1/2} \leq \left\| \sum_{i=1}^r \sum_{k \in \mathbb{Z}} c_k^i \mathbf{f}^i(\cdot - k) \right\|_2 \leq K_2 \left(\sum_{j=1}^r \sum_{k \in \mathbb{Z}} |c_k^j|^2 \right)^{1/2} \quad (6.5)$$

for any r scalar sequences $c^i = \{c_k^i\} \in \ell_2(\mathbb{Z})$, $i = 1, \dots, r$.

Now, let $\mathbf{h} = (h_1, \dots, h_r)^T \in V_0 \subset L_2^r(\mathbb{R})$. Then there exist r scalar sequences $c^i = (c_k^i \in \mathbb{R}, : k \in \mathbb{Z}) \in \ell_2(\mathbb{Z})$, $i = 1, \dots, r$, such that

$$\mathbf{h} = \sum_{k \in \mathbb{Z}} (c_k^1 \mathbf{f}^1(\cdot - k) + \dots + c_k^r \mathbf{f}^r(\cdot - k)),$$

that is

$$\begin{aligned} \begin{pmatrix} h_1 \\ \vdots \\ h_r \end{pmatrix} &= \sum_{k \in \mathbb{Z}} \left(c_k^1 \begin{pmatrix} f_1^1(\cdot - k) \\ \vdots \\ f_r^1(\cdot - k) \end{pmatrix} + \dots + c_k^r \begin{pmatrix} f_1^r(\cdot - k) \\ \vdots \\ f_r^r(\cdot - k) \end{pmatrix} \right) \\ &= \sum_{k \in \mathbb{Z}} \underbrace{\begin{pmatrix} f_1^1(\cdot - k) & \dots & f_1^r(\cdot - k) \\ \vdots & & \vdots \\ f_r^1(\cdot - k) & \dots & f_r^r(\cdot - k) \end{pmatrix}}_{\mathbf{F}(\cdot - k)} \underbrace{\begin{pmatrix} c_k^1 \\ \vdots \\ c_k^r \end{pmatrix}}_{\mathbf{c}_k}. \end{aligned}$$

Thus, any $\mathbf{h} \in V_0$ can be written as $\mathbf{h} = \mathbf{F} * \mathbf{c} = \sum_{k \in \mathbb{Z}} \mathbf{F}(\cdot - k) \mathbf{c}_k$, where $\mathbf{F} \in L_2^{r \times r}(\mathbb{R})$

and $\mathbf{c} = (\mathbf{c}_k \in \mathbb{R}^r : k \in \mathbb{Z}) \in \ell_2^r(\mathbb{Z})$. Since $\mathbf{f}^1, \dots, \mathbf{f}^r \in V_0 \subset V_1$, we have

$$\mathbf{f}^i = \sum_{k \in \mathbb{Z}} \mathbf{F}(2 \cdot -k) \mathbf{a}_k^i, \quad i = 1, \dots, r,$$

for some vector sequences $\mathbf{a}^i = (\mathbf{a}_k^i \in \mathbb{R}^r : k \in \mathbb{Z}) \in \ell_2^r(\mathbb{Z})$, $i = 1, \dots, r$. Thus, the matrix-valued function

$$\mathbf{F} = (\mathbf{f}^1 | \dots | \mathbf{f}^r) \in L_2^{r \times r}(\mathbb{R})$$

satisfies the *matrix refinement equation*

$$\mathbf{F} = \sum_{k \in \mathbb{Z}} \mathbf{F}(2 \cdot -k) \mathbf{A}_k \quad (6.6)$$

where $\mathbf{A} = (\mathbf{A}_k = (\mathbf{a}_k^1 | \dots | \mathbf{a}_k^r) \in \mathbb{R}^{r \times r} : k \in \mathbb{Z})$ is the *refinement mask*.

Now, since the subdivision scheme $S_{\mathcal{A}}$ applied to the initial sequence $\delta \mathbf{I}$ converges to a matrix function with r stable columns it follows that $\mathcal{E}_{\mathcal{A}} = \mathbb{R}^r$ so that

$$\mathbf{A}_\varepsilon(1) = \sum_{j \in \mathbb{Z}} \mathbf{A}_{\varepsilon-2j} = \mathbf{I}, \quad \varepsilon \in \{0, 1\},$$

i.e. $\mathbf{A}(1) = 2\mathbf{I}$, $\mathbf{A}(-1) = \mathbf{0}$, which is the full rank property of the mask \mathcal{A} .

The “joint stability” condition (6.5) on the function vectors \mathbf{f}^j , $j = 1, \dots, r$, implies that the function \mathbf{F} is stable in the sense that there exist two constants $K_1 \leq K_2$ such that

$$K_1 \|c\|_2 \leq \|\mathbf{F} * c\|_2 \leq K_2 \|c\|_2, \quad c \in \ell_2^r(\mathbb{Z}).$$

Summarizing the described properties, the following important result holds true:

Theorem 6.1. *The subspaces $V_j \subset L_2^r(\mathbb{R})$, $j \in \mathbb{Z}$, form an MCMRA if there exists a full rank stable matrix refinable function $\mathbf{F} \in L_2^{r \times r}(\mathbb{R})$ such that*

$$V_j = \{\mathbf{F} * c(2^j \cdot) : c \in \ell_2^r(\mathbb{Z})\}, \quad j \in \mathbb{Z}.$$

In analogy to the scalar case, we call \mathbf{F} *matrix scaling function*.

6.4 Orthogonal Multichannel Multiresolution Analysis

For many application purposes, *orthogonal MCMRAs* are the most interesting. They are equivalently characterized by the following properties:

1. The space V_0 is generated by orthonormal integer translates of the function vectors \mathbf{f}^j , $j = 1, \dots, r$.
2. The matrix scaling function \mathbf{F} is orthonormal that is

$$\langle \mathbf{F}, \mathbf{F}(\cdot + k) \rangle := \int_{\mathbb{R}} \mathbf{F}^H(x) \mathbf{F}(x + k) dx = \delta_{k0} \mathbf{I}, \quad k \in \mathbb{Z}.$$

The key for the construction of orthonormal matrix scaling functions (and the associated multichannel wavelets) is their connection with interpolatory vector subdivision schemes, as established in the following Theorem where the *canonical spectral factor* $\mathbf{A}(z)$ of a positive definite interpolatory full rank symbol $\mathbf{C}(z)$ is the symbol satisfying $\mathbf{C}(z) = \frac{1}{2} \mathbf{A}^H(z) \mathbf{A}(z)$ and $\mathbf{A}(1) = 2\mathbf{I}$, $\mathbf{A}(-1) = \mathbf{0}$.

Theorem 6.2. [2] *Let $\mathbf{C}(z)$ be the symbol of an interpolatory full rank vector subdivision scheme. Let $\mathbf{C}(z)$ be symmetric strictly positive definite for all $z \neq -1$ such that $|z| = 1$. Then the canonical spectral factor $\mathbf{A}(z)$ of $\mathbf{C}(z)$ defines a subdivision scheme $S_{\mathcal{A}}$ which converges in $L_2^{r \times r}(\mathbb{R})$ to an orthonormal matrix scaling function \mathbf{F} .*

As to the construction of such function, in [2] a simple procedure is given based on a modified *spectral factorization* of the *parahermitian matrix* $2\mathbf{C}(z)$, which takes into account the presence of some of its zeros on the unitary circle.

We can now associate a matrix wavelet to any orthogonal MCMRA in the standard way. A matrix function $\mathbf{G} \in V_1$ is called a *multichannel wavelet* for the orthogonal MCMRA if:

1. $W_j := V_{j+1} \ominus V_j = \{\mathbf{G} * c(2^j \cdot) : c \in \ell_2^r(\mathbb{Z})\}$, $j \in \mathbb{Z}$.
2. \mathbf{G} is orthonormal.

As to the existence of multichannel wavelets, we recall the following result:

Theorem 6.3. [1] *Suppose that the matrix scaling function $\mathbf{F} \in L_2^{r \times r}(\mathbb{R})$ has orthonormal integer translates. Then there exists an orthonormal wavelet $\mathbf{G} \in L_2^{r \times r}(\mathbb{R})$ satisfying the two-scale relation*

$$\mathbf{G} = \sum_{j \in \mathbb{Z}} \mathbf{F}(2 \cdot - j) \mathbf{B}_j \quad (6.7)$$

for a suitable mask $\mathcal{B} = (\mathbf{B}_j : j \in \mathbb{Z})$.

Observe that, as in the scalar case, due to the full rank properties of \mathbf{F} , the symbol of the multichannel wavelet \mathbf{G} must possess at least one factor $(z - 1)$. This is equivalent to say that the multichannel wavelet has at least one vanishing moment.

It is easy to show that the symbols of an orthonormal matrix function \mathbf{F} and of the corresponding multichannel wavelet \mathbf{G} satisfy the *orthogonality (quadrature mirror filter (QMF)) conditions*

$$\begin{aligned} \mathbf{A}^\sharp(z) \mathbf{A}(z) + \mathbf{A}^\sharp(-z) \mathbf{A}(-z) &= 4\mathbf{I}, & |z| = 1, \\ \mathbf{A}^\sharp(z) \mathbf{B}(z) + \mathbf{A}^\sharp(-z) \mathbf{B}(-z) &= \mathbf{0}, & |z| = 1, \\ \mathbf{B}^\sharp(z) \mathbf{B}(z) + \mathbf{B}^\sharp(-z) \mathbf{B}(-z) &= 4\mathbf{I}, & |z| = 1, \end{aligned}$$

where $\mathbf{A}^\sharp(z) := \mathbf{A}^T(z^{-1})$ which means $\mathbf{A}^\sharp(z) = \mathbf{A}^H(z)$ whenever $|z| = 1$.

The QMF can be written in a concise form by the condition $\mathbf{U}^\sharp(z) \mathbf{U}(z) = \mathbf{I}$ on the block matrix

$$\mathbf{U}(z) := \frac{1}{2} \begin{pmatrix} \mathbf{A}(z) & \mathbf{B}(z) \\ \mathbf{A}(-z) & \mathbf{B}(-z) \end{pmatrix}, \text{ where } \mathbf{U}^\sharp(z) := \frac{1}{2} \begin{pmatrix} \mathbf{A}^\sharp(z) & \mathbf{A}^\sharp(-z) \\ \mathbf{B}^\sharp(z) & \mathbf{B}^\sharp(-z) \end{pmatrix},$$

(which is the condition of being an unitary matrix for $|z| = 1$).

Note that, given an orthogonal symbol $\mathbf{A}(z)$, the alternating flip trick, used to construct the wavelet symbol in the scalar case, does not work in this context. Indeed, the symbol $\mathbf{B}(z) = z\mathbf{A}^\sharp(-z)$ does not verify the QMF equations unless $\mathbf{A}(z)$ and $\mathbf{A}(-z)$ commute.

On the other hand, the orthogonality conditions can be written in terms of the subsymbols $\mathbf{A}_0(z)$, $\mathbf{A}_1(z)$ and $\mathbf{B}_0(z)$, $\mathbf{B}_1(z)$. In fact, from (6.4) we see that the matrix $\mathbf{V}(z)$

$$\mathbf{V}(z) := \begin{pmatrix} \mathbf{A}_0(z^2) & \mathbf{B}_0(z^2) \\ \mathbf{A}_1(z^2) & \mathbf{B}_1(z^2) \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ \frac{1}{z} & -\frac{1}{z} \end{pmatrix} \mathbf{U}(z)$$

is such that $V^\sharp(z)V(z) = 2I$. Thus, the QMF equations take the equivalent form

$$\begin{aligned} A_0^\sharp(z^2)A_0(z^2) + A_1^\sharp(z^2)A_1(z^2) &= 2I, & |z| = 1, \\ A_0^\sharp(z^2)B_0(z^2) + A_1^\sharp(z^2)B_1(z^2) &= \mathbf{0}, & |z| = 1, \\ B_0^\sharp(z^2)B_0(z^2) + B_1^\sharp(z^2)B_1(z^2) &= 2I, & |z| = 1. \end{aligned} \quad (6.8)$$

Suppose now that we are given an orthogonal MCMRA generated by some orthonormal matrix scaling function F , with G as the corresponding multichannel wavelet. Let us consider any vector-valued function $\mathbf{h} \in L_2^r(\mathbb{R})$. From the nesting properties of the spaces $\{V_j\}$ and $\{W_j\}$, it follows that the approximation $P_\ell \mathbf{h}$ of \mathbf{h} in the space V_ℓ , $\ell \in \mathbb{Z}$, can be found in terms of the following *multichannel wavelet decomposition*:

$$P_\ell \mathbf{h} = P_{\ell-L} \mathbf{h} + Q_{\ell-1} \mathbf{h} + Q_{\ell-2} \mathbf{h} + \dots + Q_{\ell-L} \mathbf{h},$$

where $L > 0$ and $P_{\ell-L} \mathbf{h}$, $Q_{\ell-j} \mathbf{h}$, $j = 1, \dots, L$, represent the orthogonal projections of \mathbf{h} to the spaces $V_{\ell-L}$, $W_{\ell-j}$, $j = 1, \dots, L$, respectively. Analogous to the scalar case, we can derive a fast algorithm which allows to compute all the projections by means of a recursive scheme. In fact, we have that, if $P_j \mathbf{h} \in V_j$ then

$$P_j \mathbf{h} = \sum_{k \in \mathbb{Z}} F(2^j \cdot -k) \mathbf{c}_k^{(j)}$$

where

$$\mathbf{c}_k^{(j)} = \langle F(2^j \cdot -k), \mathbf{h} \rangle = \int F^H(2^j x - k) \mathbf{h}(x) dx.$$

To compute the vector coefficient sequence $\mathbf{c}^{(j-1)} \in \ell_2^r(\mathbb{Z})$ connected to the representation of \mathbf{h} in the space V_{j-1} , we make use of the refinement equation and get

$$\begin{aligned} \mathbf{c}_k^{(j-1)} &= \langle F(2^{j-1} \cdot -k), \mathbf{h} \rangle = \left\langle \sum_{n \in \mathbb{Z}} F(2(2^{j-1} \cdot -k) - n) A_n, \mathbf{h} \right\rangle \\ &= \sum_{n \in \mathbb{Z}} A_{n-2k}^T \langle F(2^j \cdot -n), \mathbf{h} \rangle = \sum_{n \in \mathbb{Z}} A_{n-2k}^T \mathbf{c}_n^j, \quad k \in \mathbb{Z}. \end{aligned}$$

In the same way, using (6.7), the wavelet vector coefficients sequence $\mathbf{d}^{(j-1)} \in \ell_2^r(\mathbb{Z})$ connected to the representation of \mathbf{h} in the space W_{j-1} is obtained as

$$\mathbf{d}_k^{(j-1)} = \langle G(2^{j-1} \cdot -k), \mathbf{h} \rangle = \sum_{n \in \mathbb{Z}} B_{n-2k}^T \mathbf{c}_n^j, \quad k \in \mathbb{Z}.$$

In summary, assuming $\ell = 0$, that is V_0 as the initial space of our representation, the *vector decomposition formula* up to the level $L > 0$ reads as

$$\mathbf{c}_k^{(j-1)} = \sum_{n \in \mathbb{Z}} A_{n-2k}^T \mathbf{c}_n^j, \quad \mathbf{d}_k^{(j-1)} = \sum_{n \in \mathbb{Z}} B_{n-2k}^T \mathbf{c}_n^j, \quad k \in \mathbb{Z}, \quad j = 0, \dots, L.$$

Conversely, given the projections $P_j \mathbf{h}$ and $Q_j \mathbf{h}$, the vector coefficient sequence $\mathbf{c}^{(j+1)}$ connected to the representation of \mathbf{h} in the space $V_{j+1} = V_j \oplus W_j$ is obtained by considering that

$$P_{j+1} \mathbf{h} = \sum_{k \in \mathbb{Z}} \mathbf{F}(2^{j+1} \cdot -k) \mathbf{c}_k^{(j+1)} \quad (6.9)$$

and, on the other hand,

$$P_{j+1} \mathbf{h} = P_j \mathbf{h} + Q_j \mathbf{h} = \sum_{n \in \mathbb{Z}} \mathbf{F}(2^j \cdot -n) \mathbf{c}_n^{(j)} + \sum_{n \in \mathbb{Z}} \mathbf{G}(2^j \cdot -n) \mathbf{d}_n^{(j)}.$$

By invoking again the refinement equation on the right-hand side of the previous expression, we have that

$$P_{j+1} \mathbf{h} = \sum_{n \in \mathbb{Z}} \sum_{k \in \mathbb{Z}} \mathbf{F}(2^{j+1} - k) \mathbf{A}_{k-2n} \mathbf{c}_n^{(j)} + \sum_{n \in \mathbb{Z}} \sum_{k \in \mathbb{Z}} \mathbf{G}(2^{j+1} - n) \mathbf{B}_{k-2n} \mathbf{d}_n^{(j)}. \quad (6.10)$$

A comparison between the two expressions (6.9) and (6.10) gives the *vector reconstruction formula*:

$$\mathbf{c}_k^{(j+1)} = \sum_{n \in \mathbb{Z}} \left(\mathbf{A}_{k-2n} \mathbf{c}_n^{(j)} + \mathbf{B}_{k-2n} \mathbf{d}_n^{(j)} \right), \quad k \in \mathbb{Z}, j = -L, \dots, 0.$$

For similar matrix wavelet decomposition and reconstruction schemes see [14].

6.5 Multichannel Wavelet Construction

The equation $\mathbf{A}_0^\sharp(z) \mathbf{A}_0(z) + \mathbf{A}_1^\sharp(z) \mathbf{A}_1(z) = 2\mathbf{I}$ is the starting point of the procedure that we propose to construct the symbol $\mathbf{B}(z)$ of the multichannel wavelet. More in detail, the procedure derives the matrix Laurent polynomials $\mathbf{B}_0(z)$ and $\mathbf{B}_1(z)$, and thus $\mathbf{B}(z) = \mathbf{B}_0(z^2) + z\mathbf{B}_1(z^2)$.

In the simple situation where $\mathbf{A}_0(z)$ and $\mathbf{A}_1(z)$ are diagonal symbols, the multichannel wavelet subsymbols $\mathbf{B}_0(z)$ and $\mathbf{B}_1(z)$ are constructed by the repeated application of a scalar procedure: for each couple $(a_0(z))_{ii}, (a_1(z))_{ii}$, $i = 1, \dots, r$, the two Laurent polynomials $(b_0(z))_{ii}, (b_1(z))_{ii}$ solution of the Bezout identity are derived (see Lemma 6.1).

In the general situation, the strategy is based on the existence of two matrix symbols $\mathbf{D}_0(z), \mathbf{D}_1(z)$ such that the positive definite Hermitian matrix

$$\mathbf{G}(z) := \begin{pmatrix} \mathbf{A}_0(z) & \mathbf{D}_0^\sharp(z) \\ \mathbf{A}_1(z) & \mathbf{D}_1^\sharp(z) \end{pmatrix} \quad (6.11)$$

satisfies

$$\det \mathbf{G}(z) = 1, \quad \mathbf{G}^\sharp(z) \mathbf{G}(z) = \begin{pmatrix} 2\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_0(z) \mathbf{D}_0^\sharp(z) + \mathbf{D}_1(z) \mathbf{D}_1^\sharp(z) \end{pmatrix}, \quad (6.12)$$

as proved in the following proposition.

Proposition 6.3. *Let $\mathbf{A}_0(z)$ and $\mathbf{A}_1(z)$ be such that $\mathbf{A}_0^\sharp(z) \mathbf{A}_0(z) + \mathbf{A}_1^\sharp(z) \mathbf{A}_1(z) = 2\mathbf{I}$. Suppose that there exist two matrix symbols $\mathbf{D}_0(z)$, $\mathbf{D}_1(z)$ such that the positive definite Hermitian matrix $\mathbf{G}(z)$ in (6.11) has the properties (6.12). Let $\mathbf{K}(z)$ be the spectral factor of $\mathbf{D}(z) := 2 \left(\mathbf{D}_0(z) \mathbf{D}_0^\sharp(z) + \mathbf{D}_1(z) \mathbf{D}_1^\sharp(z) \right)$. Then the matrix symbols*

$$\mathbf{B}_0(z) = 2\mathbf{D}_0^\sharp(z) \left(\mathbf{K}^\sharp(z) \right)^{-1}, \quad \mathbf{B}_1(z) = 2\mathbf{D}_1^\sharp(z) \left(\mathbf{K}^\sharp(z) \right)^{-1}, \quad (6.13)$$

satisfy the two last equations in (6.8).

Proof. Since $\mathbf{D}(z)$ is positive definite with determinant equal to 1, its spectral factor $\mathbf{K}(z)$ exists and satisfies $\mathbf{D}(z) = \mathbf{K}(z) \mathbf{K}^\sharp(z)$. Therefore, if we substitute (6.13) into

$$\mathbf{A}_0^\sharp(z) \mathbf{D}_0^\sharp(z) + \mathbf{A}_1^\sharp(z) \mathbf{D}_1^\sharp(z) = \mathbf{0}$$

we get

$$\mathbf{A}_0^\sharp(z^2) \mathbf{B}_0(z^2) + \mathbf{A}_1^\sharp(z^2) \mathbf{B}_1(z^2) = \mathbf{0}, \quad |z| = 1,$$

which is the second equation in (6.8). Finally, since

$$(\mathbf{K}(z))^{-1} \mathbf{D}(z) \left(\mathbf{K}^\sharp(z) \right)^{-1} = \mathbf{I},$$

by substituting (6.13) into it we end up with

$$\mathbf{B}_0^\sharp(z^2) \mathbf{B}_0(z^2) + \mathbf{B}_1^\sharp(z^2) \mathbf{B}_1(z^2) = 2\mathbf{I}, \quad |z| = 1$$

which is the last equation in (6.8). \square

For the proof of the existence of the Hermitian matrix $\mathbf{G}(z)$ we refer the reader to [11]. Its actual construction can be carried out as explained in the following sections.

6.5.1 Some Matrix Completion Results

We start by recalling a result about Bezout identities.

Theorem 6.4. [10] *For any pair of Laurent polynomials $a_1(z)$, $a_2(z)$ the Bezout identity*

$$a_1(z) b_1(z) + a_2(z) b_2(z) = 1 \quad (6.14)$$

is satisfied by a pair $b_1(z)$, $b_2(z)$ of Laurent polynomials if and only if $a_1(z)$ and $a_2(z)$ have no common zeros. Moreover, given one particular pair of solutions $b_1^*(z)$, $b_2^*(z)$ the set of all solutions of (6.14) is of the form

$$b_1(z) = b_1^*(z) + p(z)a_2(z), \quad b_2(z) = b_2^*(z) - p(z)a_1(z)$$

where $p(z)$ is any Laurent polynomial.

As to the common zeros of the subsymbols, the following result holds true.

Lemma 6.1. *The subsymbols associated with the diagonal entries of an orthogonal symbol $A(z)$, that is the subsymbols $(a_0(z))_{ii}$, $(a_1(z))_{ii}$, $i = 1, \dots, r$, have no common zeros.*

Proof. The result follows from the first equation in (6.8). In fact the existence of a common zero for $(a_0(z))_{ii}$, $(a_1(z))_{ii}$, $i = 1, \dots, r$ contradicts

$$(a_0(z^{-1}))_{ii}(a_0(z))_{ii} + (a_1(z^{-1}))_{ii}(a_1(z))_{ii} = 2, \quad i = 1, \dots, r,$$

a relation satisfied by the diagonal symbols. □

We continue with a “completion” result.

Theorem 6.5. *Let $a_i(z)$, $i = 1, \dots, n$, be Laurent polynomials with $n \geq 2$ with no common zeros. Then there exists a $n \times n$ matrix Laurent polynomial $P(z)$ whose first row is $(a_1(z), \dots, a_n(z))$ such that*

$$\det P(z) = 1, \quad z \in \mathbb{C} \setminus \{0\}.$$

Proof. The proof is by induction. Let us start with $n = 2$. For the row vector $(a_1(z), a_2(z))$ we construct the matrix $P(z) = \begin{pmatrix} a_1(z) & a_2(z) \\ -b_1(z) & b_2(z) \end{pmatrix}$ with $b_i(z)$, $i = 1, 2$ being solutions of the Bezout identity (6.14). Obviously, the matrix $P(z)$ is such that $\det P(z) = 1$, $z \in \mathbb{C} \setminus \{0\}$. Next, assuming that the theorem is true for $n - 1$, we prove it for n . Thus, given $(a_1(z), \dots, a_n(z))$ we first construct $\bar{P}(z)$, $\det \bar{P}(z) = 1$, with the first row given by $(a_1(z), \dots, a_{n-1}(z))$ then we construct the block matrix

$$P(z) := \begin{pmatrix} & a_n(z) \\ & 0 \\ & \vdots \\ \bar{P}(z) & 0 \\ 0 \ 0 \ \dots \ 0 & 1 \end{pmatrix}.$$

This completes the proof. □

With this result we are now able to describe a procedure for “completing” an $m \times n$ matrix Laurent polynomial with $m \geq n$ and $\text{rank} A(z) = m$, by constructing an

$n \times n$ matrix Laurent polynomial $\mathbf{P}(z)$, whose first m rows agree with those of $\mathbf{A}(z)$, such that

$$\det \mathbf{P}(z) = 1, \quad z \in \mathbb{C} \setminus \{0\}.$$

The existence of such polynomial has been formerly proved in [11].

We start by discussing how to write a procedure providing, for a given Laurent polynomial vector $\mathbf{a}_n(z)$ of length $n \geq 2$, a Laurent polynomial matrix $\mathbf{P}(z)$ having the first row given by $\mathbf{a}_n(z) = (a_1(z), \dots, a_n(z))$ and satisfying $\det \mathbf{P}(z) = 1$, $z \in \mathbb{C} \setminus \{0\}$.

For $a_1(z), a_2(z)$, let $[b_1(z), b_2(z)] = \text{Bezout}(a_1(z), a_2(z))$ be the procedure in Matlab-like notation providing the two Laurent polynomials $b_1(z), b_2(z)$ solution of the Bezout identity.

Next, let $[\mathbf{P}_{n \times n}(z)] = \text{Basic_completion}(\mathbf{a}_n(z))$ be the following (recursive) procedure that, taking the vector $\mathbf{a}(z)$ as input, produces the matrix $\mathbf{P}_{n \times n}(z)$ as output.

- $[\mathbf{P}_{n \times n}(z)] = \text{Basic_completion}(\mathbf{a}_n(z))$

If $n = 2$ let $[b_1(z), b_2(z)] = \text{Bezout}(a_1(z), a_2(z))$ two Laurent polynomials solution of the Bezout identity. Set $\mathbf{P}_{2 \times 2}(z) := \begin{pmatrix} a_1(z) & a_2(z) \\ -b_1(z) & b_2(z) \end{pmatrix}$;

else ($n > 2$) use $[\mathbf{P}_{(n-1) \times (n-1)}(z)] = \text{Basic_completion}(\mathbf{a}_{n-1}(z))$.

Then with $\mathbf{e}_{n-1} := (0, \dots, 1)^T \in \mathbb{R}^{n-1}$, define

$$\mathbf{P}_{n \times n}(z) := \begin{pmatrix} \mathbf{P}_{(n-1) \times (n-1)}(z) & \mathbf{a}_n(z) \\ \mathbf{0}_{1 \times (n-1)} & \mathbf{e}_{n-1} \end{pmatrix}.$$

We continue by describing the (recursive) procedure $[\mathbf{P}_{m \times m}] = \text{Completion}(\mathbf{A}_{n \times m})$. This procedure, for a given matrix Laurent polynomial $\mathbf{A}_{n \times m}(z)$ with $m > n$, constructs the Laurent polynomial matrix $\mathbf{P}(z)$ whose n first rows agree with those of $\mathbf{A}_{n \times m}$ and such that $\det \mathbf{P}(z) = 1$, $z \in \mathbb{C} \setminus \{0\}$.

- $[\mathbf{P}_{m \times m}(z)] = \text{Completion}(\mathbf{A}_{n \times m}(z))$
 1. Take $\bar{\mathbf{A}}_{(n-1) \times m}(z)$ the sub-matrix of $\mathbf{A}_{n \times m}(z)$ made of its first $(n-1)$ rows.
 2. If $n = 2$ construct $[\bar{\mathbf{P}}_{m \times m}(z)] = \text{Basic_completion}(\bar{\mathbf{A}}_{1 \times m}(z))$
else ($n > 2$) construct $[\bar{\mathbf{P}}_{m \times m}(z)] = \text{Completion}(\bar{\mathbf{A}}_{(n-1) \times m}(z))$.
 3. Compute the matrix $\mathbf{C}_{n \times m}(z) := \mathbf{A}_{n \times m}(z) \bar{\mathbf{P}}_{m \times m}^{-1}(z)$ having block structure

$$\mathbf{C}_{n \times m}(z) = \begin{pmatrix} \mathbf{I}_{(n-1) \times (n-1)} & \mathbf{0}_{(n-1) \times (r+1)} \\ \mathbf{c}_{1 \times (n-1)}(z) & \mathbf{d}_{1 \times (r+1)}(z) \end{pmatrix}$$

where $r := m - n$.

4. Construct the $(r+1) \times (r+1)$ matrix
 $[\mathbf{D}_{(r+1) \times (r+1)}(z)] = \text{Basic_completion}(\mathbf{d}_{1 \times (r+1)}(z)).$

5. Use $\mathbf{D}_{(r+1) \times (r+1)}(z)$ to construct the $m \times m$ block matrix

$$\mathbf{V}_{m \times m}(z) := \left(\begin{array}{c|c} \mathbf{I}_{(n-1) \times (n-1)}(z) & \mathbf{0}_{(n-1) \times (r+1)} \\ \hline \mathbf{c}_{1 \times (n-1)}(z) & \mathbf{D}_{(r+1) \times (r+1)}(z) \\ \mathbf{0}_{r \times (n-1)} & \end{array} \right)$$

6. Set $\mathbf{P}_{m \times m}(z) := \mathbf{V}_{m \times m}(z) \bar{\mathbf{P}}_{m \times m}(z)$.

With the help of the two previous recursive procedures, we are able to sketch the algorithm for constructing multichannel wavelets.

6.5.2 The MCW Construction Algorithm

1. From $\mathbf{A}(z) \in \ell_0^{r \times r}(\mathbb{Z})$ extract the sub-symbols $\mathbf{A}_0(z)$ and $\mathbf{A}_1(z)$.
2. If $\mathbf{A}_0(z)$ and $\mathbf{A}_1(z)$ are diagonal symbols construct $\mathbf{B}_0(z)$ and $\mathbf{B}_1(z)$ by the repeated application of the procedure
 $[(b_0(z))_{ii}, (b_1(z))_{ii}] = \text{Bezout}((a_0(z))_{ii}, (a_1(z))_{ii}), i = 1, \dots, r$
 Then go to step 10.
3. Otherwise construct the $r \times 2r$ block matrix $\tilde{\mathbf{A}}^\sharp(z) = \left(\mathbf{A}_0^\sharp(z) \mid \mathbf{A}_1^\sharp(z) \right)$.
4. Use $[\mathbf{L}(z)] = \text{Completion}(\tilde{\mathbf{A}}^\sharp(z))$ to construct the matrix

$$\mathbf{L}(z) =: \begin{pmatrix} \mathbf{A}_0^\sharp(z) & \mathbf{A}_1^\sharp(z) \\ \mathbf{C}_0(z) & \mathbf{C}_1(z) \end{pmatrix}$$

so that

$$\det \begin{pmatrix} \mathbf{A}_0(z) & \mathbf{C}_0^\sharp(z) \\ \mathbf{A}_1(z) & \mathbf{C}_1^\sharp(z) \end{pmatrix} = 1.$$

5. Compute the $r \times r$ matrix $\mathbf{R}(z) = \frac{1}{2} \left(\mathbf{A}_0^\sharp(z) \mathbf{C}_0^\sharp(z) + \mathbf{A}_1^\sharp(z) \mathbf{C}_1^\sharp(z) \right)$.
6. Compute the $r \times r$ matrices $\mathbf{D}_i(z) = -\mathbf{R}^\sharp(z) \mathbf{A}_i^\sharp(z) + \mathbf{C}_i(z), \quad i = 0, 1$.
7. Compute the $r \times r$ positive definite matrix

$$\mathbf{D}(z) = 2 \left(\mathbf{D}_0(z) \mathbf{D}_0^\sharp(z) + \mathbf{D}_1(z) \mathbf{D}_1^\sharp(z) \right).$$

8. Compute the spectral factorization of $\mathbf{D}(z)$ such that

$$\mathbf{D}(z) = \mathbf{K}(z) \mathbf{K}^\sharp(z),$$

for example using the algorithm described in [9].

9. Set $\mathbf{E}(z) = \mathbf{K}^{-1}(z)$ and construct the $r \times r$ matrices

$$\mathbf{B}_i(z) = 2\mathbf{D}_i^\sharp(z) \mathbf{E}^\sharp(z), \quad i = 0, 1.$$

10. Construct the wavelet symbol

$$\mathbf{B}(z) = \mathbf{B}_0(z^2) + z\mathbf{B}_1(z^2).$$

Two remarks are worth to be made:

- For a proof of the correctness of all steps we refer to [11, Sect. 6]. Nevertheless, since [11] deals with rank-1 matrix functions, the above construction has to be considered as an extension to the full rank case;
- The algorithm produces correct results even in case of diagonal subsymbols $\mathbf{A}_0(z)$ and $\mathbf{A}_1(z)$. Though, to end up with exactly diagonal subsymbols $\mathbf{B}_0(z)$ and $\mathbf{B}_1(z)$ we may need to multiply the result of the MCW algorithm by a suitable permutation matrix \mathbf{J} .

6.6 Numerical Examples

The aim of this section is to consider two orthonormal full rank refinable matrix functions with $r = 2$ and $r = 3$, respectively, and to use the MCW algorithm described in the Sect. 6.5.2 to construct the corresponding multichannel wavelets. In particular, while for the case $r = 2$ the positive definite full rank symbol that we start with is the one considered in [2], the case $r = 3$ is here derived from scratch.

6.6.1 A Two-Channel Example

Let $\mathbf{C}(z)$ be the symbol of the interpolatory full rank vector subdivision scheme first given in [2]

$$\mathbf{C}(z) = \begin{pmatrix} c_{11}(z) & c_{12}(z) \\ c_{12}(z^{-1}) & c_{22}(z) \end{pmatrix}$$

with

$$\begin{aligned} c_{11}(z) &= \frac{1}{2} \frac{(z+1)^2}{z} && \text{(linear B-spline symbol),} \\ c_{22}(z) &= -\frac{1}{16} \frac{(z^2 - 4z + 1)(z+1)^4}{z^3} && \text{(4-point scheme symbol),} \\ c_{12}(z) &= \lambda z(z^2 - 1)^3, \end{aligned}$$

satisfying $\mathbf{C}(1) = 2\mathbf{I}$, $\mathbf{C}(z) + \mathbf{C}(-z) = 2\mathbf{I}$. In order to assure positive definiteness, the parameter λ is taken in the interval $(-\frac{1}{32}\sqrt{3}, \frac{1}{32}\sqrt{3})$. Fixing, for example, $\lambda = \frac{1}{20}$, the symbol $\mathbf{A}(z)$ of the associated full rank orthonormal refinable function constructed with the algorithm given in [2] has elements

$$\begin{aligned} a_{11}(z) &= 1.081604742 + 0.7776021479z + 0.2165957837z^{-1} - 0.08257171989z^{-2} \\ &\quad + 0.004835070286z^{-3} + 0.0009670018466z^{-4} + 0.0009670219776z^{-5}, \\ a_{12}(z) &= -0.001208777472 + 0.001208777472z^2 + 0.001208777472z \\ &\quad - 0.001208777472z^{-1}, \\ a_{21}(z) &= -0.02117594497 - 0.005136018632z - 0.07331687185z^{-1} \\ &\quad + 0.2747671662z^{-2} - 0.0679559177z^{-3} \\ &\quad - 0.2535912212z^{-4} + 0.1464088082z^{-5}, \\ a_{22}(z) &= 0.3169890016 + 0.6830110222z^2 + 1.183011034z - 0.1830110103z^{-1}. \end{aligned}$$

The plot of the corresponding matrix scaling function is given in Fig. 6.1.

Using the MCW Algorithm we find that the symbol $\mathbf{B}(z)$ of the corresponding orthonormal two-channel wavelet has the coefficients \mathbf{B}_k , $k = -2, \dots, 7$, given in Table 6.1. The plot of the wavelet is represented in Fig. 6.2.





Fig. 6.1 Orthonormal matrix scaling function with $r = 2$

Table 6.1 Coefficients of the 2-channel orthonormal wavelet

k	B_k	k	B_k
-2	$\begin{pmatrix} 0.8140199 & 0 \\ -0.0014406 & 0 \end{pmatrix}$	3	$\begin{pmatrix} -0.0044401 & 0.0519219 \\ -0.0000359 & 0.331059 \end{pmatrix}$
-1	$\begin{pmatrix} -1.1322992 & 0 \\ 0 & 0 \end{pmatrix}$	4	$\begin{pmatrix} 0.0009821 & 0.2546764 \\ -0.0000192 & 0.1253084 \end{pmatrix}$
0	$\begin{pmatrix} 0.1915220 & 0.0032939 \\ 0.0037742 & -0.0000058 \end{pmatrix}$	5	$\begin{pmatrix} 0.0007061 & 0.1209272 \\ -0.0000047 & 0.0153323 \end{pmatrix}$
1	$\begin{pmatrix} 0.1360351 & -0.0002724 \\ -0.0018743 & 0.6524543 \end{pmatrix}$	6	$\begin{pmatrix} 0 & -0.2427089 \\ 0 & 0.0047518 \end{pmatrix}$
2	$\begin{pmatrix} -0.0065259 & -0.0133465 \\ -0.0003995 & -1.1300526 \end{pmatrix}$	7	$\begin{pmatrix} 0 & -0.1744916 \\ 0 & 0.0011525 \end{pmatrix}$

**Fig. 6.2** Orthonormal 2-channel wavelet

6.6.2 A Three-Channel Example

We now construct a 3×3 positive definite parahermitian interpolatory symbol

$$\mathbf{C}(z) = \begin{pmatrix} c_{11}(z) & c_{12}(z) & c_{13}(z) \\ c_{12}(z^{-1}) & c_{22}(z) & c_{23}(z) \\ c_{13}(z^{-1}) & c_{23}(z^{-1}) & c_{33}(z) \end{pmatrix},$$

satisfying $\mathbf{C}(1) = 2\mathbf{I}$, $\mathbf{C}(z) + \mathbf{C}(-z) = 2\mathbf{I}$. It means that, as in the previous example, the diagonal elements should be the symbols of interpolatory scalar schemes, with factor $(z+1)$ of order m_1, m_2, m_3 , respectively. The off-diagonals should certainly contain a $(z^2 - 1)$ factor and must satisfy $c_{ij}(z) = -c_{ij}(-z)$ (see [4]). We take on the diagonal the Deslaurier–Dubuc filters with $m_1 = 2, m_2 = 4, m_3 = 2$, that is:

$$c_{11}(z) = \frac{1}{2} \frac{(z+1)^2}{z}, \quad c_{22}(z) = -\frac{1}{16} \frac{(z^2 - 4z + 1)(z+1)^4}{z^3}, \quad c_{33}(z) = c_{11}(z).$$

Since $(z+1)^2$ is a common factor, repeating the consideration done in [2], we require the following symbols on the off-diagonal

$$c_{12}(z) = \lambda_1 z(z^2 - 1)^3, \quad c_{13}(z) = \lambda_2 z(z^2 - 1)^4, \quad c_{23}(z) = \lambda_3 z(z^2 - 1)^4.$$

Table 6.2 Coefficients of the 3-channel orthonormal matrix scaling function

k	A_k	k	A_k
-8	$\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0.04 & 0.03125 & 0 \end{pmatrix}$	-2	$\begin{pmatrix} -0.0715410 & 0.0216501 & 0 \\ 0.2697914 & -0.0206268 & 0 \\ -0.04 & -0.03125 & 0 \end{pmatrix}$
-7	$\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -0.04 & -0.03125 & 0 \end{pmatrix}$	-1	$\begin{pmatrix} 0.2874359 & 0.0253824 & 0 \\ -0.0756559 & -0.1655344 & 0 \\ 0.04 & 0.03125 & 0 \end{pmatrix}$
-6	$\begin{pmatrix} -0.0022120 & -0.0017281 & 0 \\ 0.0000731 & 0.0000571 & 0 \\ -0.12 & -0.09375 & 0 \end{pmatrix}$	0	$\begin{pmatrix} 1.0740925 & -0.0422890 & 0 \\ -0.0129914 & 0.3501563 & 0 \\ 0 & 0 & 1 \end{pmatrix}$
-5	$\begin{pmatrix} 0.0087982 & 0.0060117 & 0 \\ 0.1481828 & -0.0017879 & 0 \\ 0.12 & 0.09375 & 0 \end{pmatrix}$	1	$\begin{pmatrix} 0.7231598 & -0.0073621 & 0 \\ -0.0053020 & 1.1683135 & 0 \\ 0 & 0 & 1 \end{pmatrix}$
-4	$\begin{pmatrix} -0.0003396 & 0.0004183 & 0 \\ -0.2568731 & 0.0059962 & 0 \\ 0.12 & 0.09375 & 0 \end{pmatrix}$	2	$\begin{pmatrix} 0 & 0.0219488 & 0 \\ 0 & 0.6644173 & 0 \\ 0 & 0 & 0 \end{pmatrix}$
-3	$\begin{pmatrix} -0.0193938 & -0.0240319 & 0 \\ -0.0672249 & -0.0009911 & 0 \\ -0.12 & -0.09375 & 0 \end{pmatrix}$		

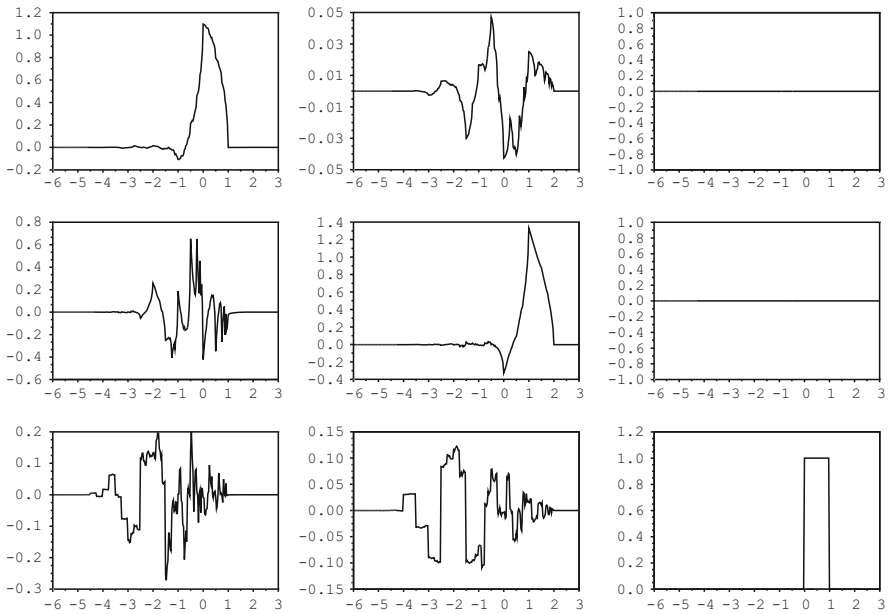


Fig. 6.3 Orthonormal matrix scaling function with $r = 3$

The following parameter choice

$$\lambda_1 = 1/20, \lambda_2 = 1/50, \lambda_3 = 1/64$$

gives positive definiteness. In this case, the canonical spectral factor $A(z)$, symbol of the orthonormal refinable function F , has the coefficients given in Table 6.2. The plot of the corresponding scaling function is represented in Fig. 6.3.

The MCW algorithm applied to this symbol, gives the orthonormal wavelet whose symbol coefficients are given in Table 6.3 and whose plot is represented in Fig. 6.4.

6.7 Conclusions

This paper discusses a way to construct orthonormal multichannel wavelets from convergent full rank orthogonal symbols. As to our knowledge, the explicit algorithm for multichannel wavelet construction (the MCW construction algorithm) here derived, based on the strategy for rank-1 filters given in [11], is the first algorithm proposed so far in the literature.

Table 6.3 Coefficients of the 3-channel orthonormal matrix wavelet

k	B_k	k	B_k
-2	$\begin{pmatrix} -0.0031659 & 0 & 0 \\ 0.0001046 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$	5	$\begin{pmatrix} 0.1684248 & 0.0212803 & 0.0595696 \\ -0.0017438 & 0.0229928 & 0.1406681 \\ -0.0001092 & -0.0004847 & -0.0160288 \end{pmatrix}$
-1	$\begin{pmatrix} 0 & 0 & 0 \\ -0.6416072 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$	6	$\begin{pmatrix} -0.0017816 & -0.0078392 & -0.1305600 \\ -0.0009269 & 0.0105339 & 0.0451634 \\ 0 & 0 & -0.0026508 \end{pmatrix}$
0	$\begin{pmatrix} 0.0297628 & -0.7705220 & -0.0001305 \\ 1.1269644 & 0.0254539 & 0 \\ 0.0367180 & 0.0169580 & -0.9720037 \end{pmatrix}$	7	$\begin{pmatrix} -0.0012167 & -0.0056664 & -0.0773112 \\ -0.0015390 & -0.0062807 & -0.1222042 \\ 0 & 0 & 0.0026508 \end{pmatrix}$
1	$\begin{pmatrix} -0.0940031 & 1.1447719 & 0 \\ -0.3500526 & -0.0169113 & -0.0264419 \\ -0.0367180 & -0.0169580 & 0.9720037 \end{pmatrix}$	8	$\begin{pmatrix} -0.0000291 & -0.0001188 & 0.0406917 \\ -0.0008803 & -0.0035956 & -0.0543552 \\ 0 & 0 & 0 \end{pmatrix}$
2	$\begin{pmatrix} -0.2583832 & -0.2560664 & -0.0751903 \\ -0.1287475 & -0.0115762 & 0.0489688 \\ -0.0096495 & -0.0033517 & -0.0409056 \end{pmatrix}$	9	$\begin{pmatrix} 0 & 0 & 0.0295229 \\ 0 & 0 & 0.0373432 \\ 0 & 0 & 0 \end{pmatrix}$
3	$\begin{pmatrix} -0.0908757 & -0.1606346 & 0.0018325 \\ -0.0045318 & -0.0170932 & -0.0567853 \\ 0.0096495 & 0.0033517 & 0.0409056 \end{pmatrix}$	10	$\begin{pmatrix} 0 & 0 & 0.0007057 \\ 0 & 0 & 0.0213611 \\ 0 & 0 & 0 \end{pmatrix}$
4	$\begin{pmatrix} 0.2512676 & 0.0347952 & 0.1508696 \\ 0.0029602 & -0.0035235 & -0.0337223 \\ 0.0001092 & 0.0004847 & 0.0160288 \end{pmatrix}$		

As already pointed out, multichannel wavelets provide an effective tool for the analysis of multichannel signals, that is vector-valued signals whose components come from different sources with possible intrinsic correlations, for example, seismic waves, brain activity (EEG/MEG) data, financial time series, color images, etc. Many multichannel signals exhibit a high correlation which can be revealed and exploited by a MCW analysis, with filters suitably tailored to the specific data and application. Future investigations include the construction of “data-adapted” MCW bases and their application in many problems where vector-valued signals have to be processed for compression, denoising, etc.

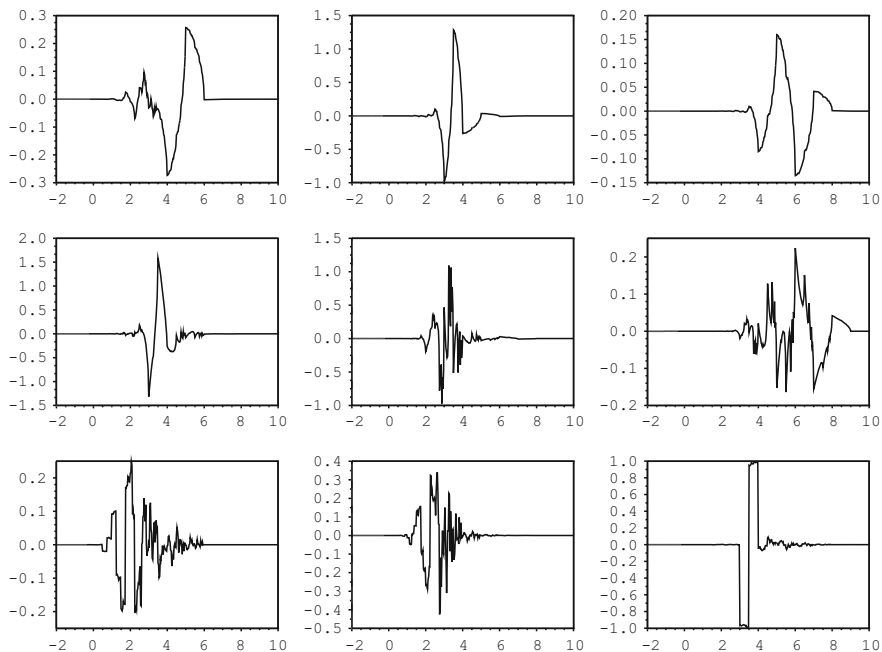


Fig. 6.4 Orthonormal 3-channel wavelet

References

1. S. Bacchelli, M. Cotronei, T. Sauer, Wavelets for multichannel signals, *Adv. Appl. Math.*, **29**, 581–598 (2002)
2. C. Conti, M. Cotronei, T. Sauer, Full rank interpolatory subdivision schemes: Kronecker, filters and multiresolution, *J. Comput. Appl. Math.*, **233** (7), 1649–1659 (2010)
3. C. Conti, M. Cotronei, T. Sauer, Full rank positive matrix symbols: interpolation and orthogonality, *BIT*, **48**, 5–27 (2008)
4. C. Conti, M. Cotronei, T. Sauer, Interpolatory vector subdivision schemes, in: A. Cohen, J. L. Merrien, L. L. Schumaker (eds.), *Curves and Surfaces*, Avignon 2006, Nashboro Press, (2007)
5. M. Cotronei, T. Sauer, Full rank filters and polynomial reproduction, *Comm. Pure Appl. Anal.*, **6**, 667–687 (2007)
6. J.E. Fowler, L. Hua, Wavelet Transforms for Vector Fields Using Omnidirectionally Balanced Multiwavelets, *IEEE Transactions on Signal Processing*, **50**, 3018–3027 (2002)
7. J. S. Geronimo, D. P. Hardin, P. R. Massopust, Fractal functions and wavelet expansions based on several scaling functions, *J. Approx. Theory*, **78** (3), 373–401 (1994)
8. F. Keinert, *Wavelets and Multiwavelets*, Chapman & Hall/CRC, (2004)
9. J. Ježek and V. Kučera, Efficient algorithm for matrix spectral factorization, *Automatica*, **21**, 663–669 (1985)
10. Lawton, W. M.; Micchelli, C. A., Bezout identities with inequality constraints. *Vietnam J. Math.* **28** (2), 97–126 (2000)
11. C. A. Micchelli, T. Sauer, Regularity of multiwavelets, *Adv. Comput. Math.*, **7** (4), 455–545 (1997)
12. G. Strang, V. Strela, Short wavelets and matrix dilation equations, *IEEE Trans. Signal Process.*, **43** (1), 108–115 (1995)

13. X.G. Xia, Orthonormal matrix valued wavelets and matrix Karhunen–Loève expansion. In: A. Aldroubi, E. B. Lin (eds.) *Wavelets, multiwavelets, and their applications*, Contemporary Mathematics, **216**, 159–175. Providence, RI: American Mathematical Society (1998)
14. X.G. Xia, B. Suter, Vector-valued wavelets and vector filter banks, *IEEE Trans. Signal Process.*, **44**, 508–518 (1996)
15. A. T. Walden, A. Serroukh, Wavelet analysis of matrix-valued time-series, *Proc. R. Soc. Lond. A*, **458**, 157–179 (2002)

Chapter 7

Unitary Systems and Bessel Generator Multipliers

Deguang Han and David R. Larson

Abstract A Bessel generator multiplier for a unitary system is a bounded linear operator that sends Bessel generator vectors to Bessel generator vectors. We characterize some special (but useful) Bessel generator multipliers for the unitary systems that are ordered products of two unitary groups. These include the wavelet and the Gabor unitary systems. We also provide a detailed exposition of some of the history leading up to this work.

7.1 Introduction

Orthonormal wavelets can be regarded as complete wandering vectors for a special system of unitary operators acting on a separable infinite dimensional Hilbert space. In the standard wavelet (and frame–wavelet) theory, the system is an ordered product $\{D^n T^k : n \in \mathbb{Z}, k \in \mathbb{Z}^d\}$ of two unitary groups $\{D^n\}$ and $\{T^k\}$ acting on the Hilbert space $L^2(\mathbb{R}^d)$ called the dilation group and the translation group, respectively, which intertwine in their action on H in a particular way. The standard Gabor frame theory has a similar operator model, with the difference being in the way the groups intertwine. In the standard wavelet theory, both groups are abelian, with the dilation group being cyclic with generator D corresponding to dilation of functions by a $d \times d$ invertible real matrix on $L^2(\mathbb{R}^d)$, and with the translation group having d generators corresponding to translation of functions by the elements of a basis of \mathbb{R}^d . In the standard Gabor theory, the groups are the modulation group and the translation group. In the Gabor theory, the groups commute modulo scalars, but in the wavelet theory they do not. There have been several generalizations of these theories by introducing other natural candidates for the dilation, translation, and modulation groups, and considering systems which are ordered products of more than two groups.

D.R. Larson (✉)

Department of Mathematics, Texas A & M University, College Station, TX, USA

e-mail: larson@math.tamu.edu

If we let $\mathcal{U}_{D,T}$ denote the wavelet unitary system with D the dilation unitary group and T the translation unitary group, then an orthonormal wavelet (respectively Riesz wavelet) can be thought as an orthonormal (respectively Riesz) basis generator vector for $\mathcal{U}_{D,T}$ and a frame wavelet can be viewed as a frame generator vector for $\mathcal{U}_{D,T}$. These have been studied in many previous papers in the literature. We will call a vector x a Bessel wavelet if the action of $\mathcal{U}_{D,T}$ on x generates a Bessel sequence; i.e. if x is a Bessel sequence generator for $\mathcal{U}_{D,T}$. It turns out that one can gain much additional perspective on wavelets and frame wavelets if one views them as special cases of Bessel wavelets. One reason for this is that the set of Bessel wavelets is a linear space, unlike the sets of wavelets and frame wavelets. When one puts a certain natural norm on the space of Bessel wavelets, *the Bessel bound norm*, in addition to the usual norm as a vector in Hilbert space, it becomes isomorphic to an *operator space* in the usual sense of the term. Thus, it admits some investigation using techniques of modern operator algebra. Some new results for wavelets and frame wavelets, and some new proofs of previously known results, can then be obtained because of their role as very special elements of this Bessel wavelet operator space. There is nothing sacred about wavelet theory in this approach. Gabor–Bessel generators are defined in the same way for the Gabor unitary systems, and they too are operator spaces in a natural way. More general unitary systems (in its greatest generality a unitary system is simply defined as a discrete collection of unitary operators that contains the identity operator I) can also be associated with operator spaces in this way.

A Bessel generator multiplier for a unitary system is a bounded linear operator that sends Bessel generator vectors to Bessel generator vectors. In this article, we will characterize some special (but useful) Bessel generator multipliers for the unitary systems that are ordered products of two unitary groups. These include the wavelet and the Gabor unitary systems discussed above. These results are new and have not been published elsewhere. We first defined Bessel generator multipliers in the earlier paper [22], where we used them to give a new proof of a theorem of Bownik [2] showing that the set of wavelet frames for a fixed wavelet unitary system $\mathcal{U}_{D,T}$ are, indeed, dense in the underlying Hilbert space $L^2(\mathbb{R}^d)$. Here, D is the dilation group corresponding to dilation by an expansive matrix. This answered a question of the second author (c.f. [31, 32]). More precisely, we first worked out the new proof of Bownik’s result and then we realized that it could be made more crystal clear by introducing the concept of a Bessel generator multiplier. In this article, we first give an exposition of our work in [22], and then we prove our new results.

We recall and introduce some notations and definitions that will be needed in this paper.

A *frame* for a Hilbert space H is a sequence of vectors $\{f_j\}_{j \in \mathbb{J}}$ in H such that there exist positive constants A and B such that

$$A\|f\|^2 \leq \sum_{j \in \mathbb{J}} |\langle f, f_j \rangle|^2 \leq B\|f\|^2$$

holds for every $f \in H$, where \mathbb{J} is a countable index set. We call A a lower frame bound and B an upper frame bound. The optimal lower and upper bounds are called

the *frame bounds*. A *tight frame* refers to a frame when both upper and lower frame bounds are the same, and a *Parseval frame* is a frame whose frame bounds are one. If we only require the right hand inequality in the frame definition, then $\{f_n\}$ is called a *Bessel sequence* and the optimal B is called the *Bessel bound*. The *analysis operator* Θ for a Bessel sequence is defined by

$$\Theta(f) = \sum_{j \in \mathbb{J}} \langle f, f_j \rangle e_j, \quad f \in H,$$

where $\{e_j\}$ is the standard orthonormal basis for the $\ell^2(\mathbb{J})$ -sequence space. It is easy to verify that $\Theta^*e_j = f_j$, $j \in \mathbb{J}$, and that $\Theta^*\Theta = \sum_j f_j \otimes f_j$, where $f \otimes g$ is the elementary tensor rank-one operator defined by $(f \otimes g)(h) = \langle h, g \rangle f$. Moreover, a sequence $\{f_j\}$ is a frame (respectively a Bessel sequence) if and only if $\sum_{j \in \mathbb{J}} f_j \otimes f_j$ is weakly convergent to a bounded invertible operator (respectively bounded operator) on H . Two Bessel sequences $\{f_j\}$ and $\{g_g\}$ are called *strongly disjoint* or *orthogonal* if the range spaces of their analysis operators are orthogonal, i.e.,

$$\sum_{j \in \mathbb{J}} \langle f, f_j \rangle \langle g_j, g \rangle = 0, \quad \forall f, g \in H.$$

There is an interesting characterization of orthogonal Parseval frames in terms of the frame property for orthogonal direct sum spaces:

Proposition 7.1. [20] *Let $\{f_j\}$ and $\{g_g\}$ be two Parseval frames for H and K , respectively. Then $\{f_j\}$ and $\{g_g\}$ are strongly disjoint if and only if $\{f_j \oplus g_g\}$ is a Parseval frame for the orthogonal direct sum space $H \oplus K$.*

The investigation of frames with special structures has been the main focus in the literature of frame theory. Included are the wavelet frames, the Gabor (or Weyl-Heisenberg) frames, and the frames obtained by the action of discrete group representations. All of these frames can be obtained by applying a collection of unitary operators, called a *unitary system*, to a *generating vector* in the representing Hilbert space. Such a vector $x \in H$ is called a (*complete*) *wandering vector* (resp *a frame generator vector, tight frame generator or a Bessel generator vector* for a unitary system \mathcal{U} if $\mathcal{U}x := \{Ux\}_{U \in \mathcal{U}}$ (as a sequence indexed by \mathcal{U}) is an orthonormal basis for H (respectively a frame, tight frame or a Bessel sequence). If a bounded linear operator A on H maps every wandering vector of \mathcal{U} to a wandering vector of U , then it is called a *wandering vector multiplier*. wandering vector multipliers play a significant role in studying the topological properties of the set of all wavelets. For example, a complete characterization of the wavelet multipliers A of the form $\hat{A} = M_h$ was obtained in [39] and was used to show that the set of all MRA wavelets is path-connected. Here M_h denotes the multiplication unitary operator associated with the unimodular function symbol h and $\hat{A} = \mathcal{F}A\mathcal{F}^{-1}$, where \mathcal{F} is the Fourier transform. It was shown in [21, 27] that when the unitary system is a group, then the set of all wandering vector multipliers form a group. A complete characterization of

the wandering vector multipliers for abelian and for ICC groups is also given in [21]. However, in general, it is difficult to characterize the wandering vector multipliers even when the unitary system \mathcal{U} has a relatively simple structure.

If a bounded linear operator A on H maps every frame generator vector for \mathcal{U} to a frame generator vector of U , then it is called a *frame generator multiplier*. Similarly, if A maps the set of Bessel generator vectors for U into itself it is called a *Bessel generator multiplier*.

As mentioned above, the set of wandering vector multipliers can have a complex structure even when the unitary system is relatively simple. The set of frame generator multipliers likewise has a complex structure. The wandering vector multipliers form a unitary semigroup, and in some special cases this is a group. But proofs can be difficult [21, 27]. The frame generator multipliers naturally form a multiplicative semigroup, but even in relatively simple cases not much is known about the structure.

On the other hand, it is easily seen that the set of Bessel generator vectors for a unitary system \mathcal{U} , denoted by $\mathcal{B}_{\mathcal{U}}$, is closed under the taking of linear combinations so is a linear space. It follows that the set of Bessel generator multipliers for \mathcal{U} is an algebra of operators in $B(H)$. It is not generally norm closed, and it will not generally be self-adjoint. We will show that the Bessel bound is a complete norm on the space of Bessel generator vectors, making it a Banach space. In fact it is an operator space, which can be seen in more than one natural way. The operator space structure resulting from these different natural representations are equivalent even though the constructions in the representations are different. We will show that a Bessel generator multiplier A (which is a bounded linear operator on H) is also bounded as a linear map from $(\mathcal{B}_{\mathcal{U}}$, the Bessel bound norm) into itself. Thus the algebra of Bessel generator multipliers carries two norms: the usual operator norm from $B(H)$, and the operator norm from the Banach space $(\mathcal{B}_{\mathcal{U}}$, the Bessel bound norm) into itself. These norms are frequently different.

A Bessel generator multiplier A does not usually commute with the action of the unitary system, so while it induces a bounded endomorphism of the operator space $\mathcal{B}_{\mathcal{U}}$, it is not usually a simple endomorphism of the type induced by a multiplication operator on a Banach algebra. If x is in $\mathcal{B}_{\mathcal{U}}$ and $y = Ax$ then the Bessel sequence $\mathcal{U}y$ generated by the action of the unitary system \mathcal{U} on y is generally *not* the sequence of vectors indexed by \mathcal{U} obtained by multiplying the sequence $\mathcal{U}x$ on the left by A . The action of A on y is in a sense the action of A on x *twisted* by the action of the unitary group on A . This twisting is what makes it interesting, and makes the structure of the generator multiplier algebra rather rich. There will be no twisting if and only if A actually commutes with the unitary system \mathcal{U} . We will give some examples to illustrate this. There are some interesting open questions involving the operator space structures of the space of Bessel generator vectors and algebra of Bessel generator multipliers. We will present some results in this direction, and pose some open questions in this direction.

This article was written for inclusion in a special volume of proceedings of the conference *Twenty Years of Wavelets*, Chicago, May 2009. The second author thanks the organizers for inviting him to be a plenary speaker.

7.2 Density and Connectivity of Wavelet Frames

As we mentioned in the introduction, one of the motivations for studying Bessel vector multipliers comes from its connection with the density and connectivity problems for wavelet frames. In a recent paper [22], we showed that the algebra of all Bessel generator multipliers for a wavelet system carries a rich algebraic and topological structure that can be used to give a solution to the path-connectivity problem and the density problem for wavelet frames. Bownik solved these problems in [2]. We gave an alternate solution which is more transparent in some ways. Since the wavelet Bessel generator multipliers were proven useful, and with seemingly *math-e-magical* properties, in the study of certain properties of wavelets, we were led to consider a more complete investigation of their properties, not only for wavelet systems but for other natural unitary systems.

The question of whether the set of all Riesz wavelets for the dyadic wavelet system on $L^2(\mathbb{R})$ is a norm-dense path-wise connected subset of $L^2(\mathbb{R})$ was asked in [29] by the second author. This question is related to the problem that was posed by Guido Weiss and his group [25, 26] and independently by Dai and Larson ([8], problem 1) of whether the set of all orthonormal dyadic wavelets is connected. Neither conjecture has been settled to date, although it has been shown [39] that the set of *MRA* orthonormal dyadic wavelets is connected. Shortly after this, the two authors of this paper developed an operator-theoretic approach to frame theory in [20], and the same problems (density and connectivity) were posed for wavelet frames (also called frame wavelets or framelets). While this density problem was not posed formally in [20] it was alluded to in that memoir, and also in the semi-expository papers [30] and [32], and it was finally posed formally (along with some other related problems) in [31]. However, it was pointed out in [20] that for the Gabor-frame case, and for the group-frame case (i.e. frame sequences generated by the action of a unitary group on a single generator vector), both the density and the connectivity problems have positive solutions. This follows from the parametrization theorem for frame vectors in terms of the set of all invertible operators from the von Neumann algebra generated by the unitary system. It was later proven that the parametrization result presented in [20] is also valid and very useful for the more general *projective* unitary group representations (cf. [10, 12, 13, 18]). These systems were also called *group-like* systems, and they include the Gabor systems as special cases. The frame density and connectivity problems also have positive solutions for these systems. The general connectivity and density problems for the set of wavelet frames remained open until recently Bownik settled both problems affirmatively.

Theorem 7.1. [2] *The set of all wavelet frames is dense and path-connected in $L^2(\mathbb{R}^d)$.*

Bownik in fact proved the density of wavelet frames not only in the Hilbert space norm but also in another natural log-modular norm associated with wavelet theory (see [36] and [17] for the definition of this norm). While the general connectivity and

density problems are settled for the case of wavelet **frames**, they are still open for the case of wavelets. We point out that before Bownik's path connectivity result, there had been steady progress on the connectivity problem for special classes of wavelets and wavelet frames, including the MSF (s-elementary) and MRA (multiresolution analysis) wavelets and their analogues for wavelet frames and especially Parseval wavelets. The interested reader should refer to (c.f. [4, 6, 7, 14–16, 23, 33, 34, 37, 39]) for details and exposition of these results. Bownik's result does not overlap with the results for special classes of wavelet frames since it is usually more difficult and challenging to construct or prove the existence of continuous paths in a subset of wavelet frames.

In [22], we re-examined the density and connectedness results of Bownik and observed that the essential ingredient that *makes the proofs work* is the *frame-orthogonality (or strong-disjointness)*, a concept that arose simultaneously and independently to Balan [1] and the authors some time ago [20]. This is a natural geometric concept in frame theory which was formally introduced and studied in [1, 20], and it has proven useful for developing the theory of frames and its applications. While Bownik does not explicitly use term frame-orthogonality in his argument, it is clear that the essential reason the argument works seems to involve the orthogonality concept. This was made *crystal-clear* in the argument of [22] where a key observation is that if a wavelet Bessel function f is strongly disjoint with a wavelet frame ψ , then $f + t\psi$ is a wavelet frame for all $t \neq 0$ and hence f is the limit of a sequence of wavelet frames. This allows us to obtain a new type of approximation result for wavelet Bessel sequence generators.

Lemma 7.1. [22] *Assume that ψ is wavelet frame and g is Bessel. If ψ and g are strongly disjoint, then $\varphi = \psi + g$ is also a wavelet frame. Moreover, if ψ has upper and lower frame bounds A and B , respectively, and g has a Bessel bound C , then φ has a lower frame bound A , and upper frame bound $B + C$.*

With the help of Lemma 7.1, we obtained the following stronger version of Bownik's density theorem.

Theorem 7.2. [22] (i) *Every function $f \in L^2(\mathbb{R})$ is a limit of an asymptotically tight sequence of wavelet frames, i.e., there exists a sequence of wavelet frames $\{\psi_n\}$ such that $\lim_{n \rightarrow \infty} \frac{B_n}{A_n} = 1$ and $\lim_{n \rightarrow \infty} \psi_n = f$, where B_n and A_n are the upper and lower frame bounds of ψ_n .*

(ii) *Assume that ψ is a wavelet frame and $f \in L^2(\mathbb{R}) \cap L^\infty(\mathbb{R})$. Then there exists a continuous path ψ_t such that $\psi_0 = \psi$, $\psi_1 = f$ and ψ_t is a wavelet frame for every $0 \leq t < 1$. In particular, this also implies that the set of all wavelet frames is path-connected and is dense in $L^2(\mathbb{R})$.*

All these results still hold for wavelet frames with expansive dilation matrices and integer lattice \mathbb{Z}^d -translations in $L^2(\mathbb{R}^d)$.

Theorem 7.2 (i) seems indicate that the set of all tight wavelet frames might be dense in $L^2(\mathbb{R})$. This was proven to be false by Bownik in [3]. In fact he proved that

if a function $\psi \in L^2(\mathbb{R})$ is in the closure of the set of all tight frames, then it satisfies one of the familiar equations characterizing Parseval wavelet frames, i.e.,

$$\sum_{j=0}^{\infty} \hat{\psi}(2^j) \overline{\hat{\psi}(2^j(\xi + q))} = 0$$

holds for all a.e. $\xi \in \mathbb{R}$ and $q \in 2\mathbb{Z} + 1$. Consequently, the set of all tight wavelet frames can not be dense in $L^2(\mathbb{R})$. Similar results hold for tight frames induced by group representations (c.f. Example 2 in [22]).

7.3 The Bessel Vector Spaces

Let \mathcal{U} be a unitary system on H and $\mathcal{B}_{\mathcal{U}}$ be the set of all the Bessel vectors for \mathcal{U} . It is easy from the definition of Bessel sequences to prove that $\mathcal{B}_{\mathcal{U}}$ is a linear subspace of H . Since the mapping from $\mathcal{B}_{\mathcal{U}}$ to $B(H, \ell^2(\mathcal{U}))$ defined by $x \rightarrow \Theta_x$ is linear and injective, we can identify $\mathcal{B}_{\mathcal{U}}$ with the space of all the analysis operators $\{\Theta_x : x \in \mathcal{B}_{\mathcal{U}}\}$ which is also a linear space. So we have two natural norms on $\mathcal{B}_{\mathcal{U}}$: The Hilbert space norm $\|x\|$ and the operator norm $\|\Theta_x\|$. Let $\{\chi_U : U \in \mathcal{U}\}$ be the standard orthonormal basis for $\ell^2(\mathcal{U})$. Then $\Theta_x^* \chi_U = Ux$ for all $U \in \mathcal{U}$. So we always have

$$\|x\| = \|\Theta_x^* \chi_I\| \leq \|\Theta_x^*\| = \|\Theta_x\|,$$

where I is the identity operator.

In general, $\mathcal{B}_{\mathcal{U}}$ is not necessarily closed with the Hilbert space norm. But the following proposition shows that it is a Banach space with the operator norm.

Proposition 7.2. *The linear space $\mathcal{B}_{\mathcal{U}}$ is a Banach space with the operator norm, and $\{\Theta_x : x \in \mathcal{B}_{\mathcal{U}}\}$ is a closed subspace of $B(H, \ell^2(\mathcal{U}))$.*

Proof. Let $\{\Theta_{x_n}\}$ be a Cauchy sequence. Then there exists a bounded linear operator $T \in B(H, \ell^2(\mathcal{U}))$ such that $\|\Theta_{x_n} - T\| \rightarrow 0$. Thus we have $\|\Theta_{x_n}^* - T^*\| \rightarrow 0$, in particular we have

$$\|x_n - T^* \chi_I\| = \|\Theta_{x_n}^* \chi_I - T^* \chi_I\| \rightarrow 0.$$

So it suffices to check that x is Bessel and $T = \Theta_x$ with $x = T^* \chi_I$. In fact, for any $y \in H$ and any finite subset Λ of \mathcal{U} , we have

$$\begin{aligned} \sum_{U \in \Lambda} |\langle y, Ux \rangle|^2 &= \lim_{n \rightarrow \infty} \sum_{U \in \Lambda} |\langle y, Ux_n \rangle|^2 \\ &\leq \limsup_{n \rightarrow \infty} \sum_{U \in \mathcal{U}} |\langle y, Ux_n \rangle|^2 \\ &= \limsup_{n \rightarrow \infty} \|\Theta_{x_n}(y)\|^2 \\ &\leq \limsup_{n \rightarrow \infty} \|\Theta_{x_n}\|^2 \|y\|^2 \leq K \|y\|^2, \end{aligned}$$

where $K = \sup\{|\Theta_{x_n}|^2 : n \in \mathbb{N}\} < \infty$. Thus, x is Bessel. Finally, we show that $T = \Theta_x$. In fact, for any $U \in \mathcal{U}$ we have $\lim_{n \rightarrow \infty} \|\Theta_{x_n}^* \chi_U - T^* \chi_U\| = 0$. On the other hand, since $\Theta_{x_n}^* \chi_U = Ux_n$ and $\Theta_x^* \chi_U = Ux$, we also have

$$\lim_{n \rightarrow \infty} \|\Theta_{x_n}^* \chi_U - \Theta_x^* \chi_U\| = \lim_{n \rightarrow \infty} \|Ux_n - Ux\| = \lim_{n \rightarrow \infty} \|x_n - x\| = 0.$$

Thus, $T^* \chi_U = \Theta_x^* \chi_U$ for all $U \in \mathcal{U}$. Therefore, we get $\Theta_x = T$, as claimed. \square

Note that $\|\Theta_x\|^2$ is the Bessel bound for x . So for the rest of the paper we will use $\|x\|_o$ to denote $\|\Theta_x\|$, and sometimes we refer it as the Bessel bound norm of x . Thus, Proposition 7.2 states that the space $(\mathcal{B}_{\mathcal{U}}, \|\cdot\|_o)$ is a Banach space. It is also an operator space with the set of matrix norms induced from the set of matrix norms of the operator space $\{\Theta_x : x \in \mathcal{B}_{\pi}\}$ considered as a closed subspace of $B(H, \ell^2(\mathcal{U}))$. We refer to [35] for more basic information about operator spaces.

In the case that \mathcal{U} has a complete wandering vector ψ (e.g. the wavelet unitary system $\mathcal{U}_{D,T}$), we use $\mathcal{C}_{\psi}(\mathcal{U})$ to denote the *local commutant* at ψ , which is defined by

$$\mathcal{C}_{\psi}(\mathcal{U}) = \{T \in B(H) : TU\psi = UT\psi, \forall U \in \mathcal{U}\}.$$

The we can also identify $\mathcal{B}_{\mathcal{U}}$ with $\mathcal{C}_{\psi}(\mathcal{U})$ by the following lemma:

Lemma 7.2. *Assume that ψ is a (complete) wandering vector for \mathcal{U} . Then η is a Bessel vector for \mathcal{U} if and only if there exists $B \in \mathcal{C}_{\psi}(\mathcal{U})$ such that $\eta = B\psi$. Moreover, this operator B is unique, and $\|\eta\|_o = \|B\|$.*

For self containment we give a short proof.

Proof. For each $\eta \in \mathcal{B}_{\mathcal{U}}$, we define B_{η} by $B_{\eta}U\psi = U\eta$ for every $U \in \mathcal{U}$. Then we have that $B_{\eta} \in \mathcal{C}_{\psi}(\mathcal{U})$ and $B_{\eta}\psi = \eta$. Since $\{U\eta\}_{U \in \mathcal{U}}$ is an orthonormal basis, we have that $\|B_{\eta}\| = \|\Theta_{\eta}^*\| = \|\Theta_{\eta}\| = \|\eta\|_o$. For uniqueness, assume that $B' \in \mathcal{C}_{\psi}(\mathcal{U})$ such that $B'\psi = \eta$. Then we have for every $U \in \mathcal{U}$, $B'U\psi = UB'\psi = UB_{\eta}\psi = B_{\eta}U\psi$. Thus, $B' = B_{\eta}$, as claimed. \square

7.4 Bessel Generator Multipliers

7.4.1 Examples of Bessel Generator Multipliers

A bounded linear operator A on H is a Bessel generator multiplier for \mathcal{U} if it maps every Bessel vector of \mathcal{U} to a Bessel vector of \mathcal{U} . In other words, a Bessel generator multiplier is a bounded linear operator A such that $\mathcal{B}_{\mathcal{U}}$ is *invariant* under A . Here are some easy examples of Bessel generator multipliers for a general unitary system.

Example 7.1. Let \mathcal{U} be a unitary system on H .

- (i) If A commutes with \mathcal{U} , then A is Bessel generator multiplier.
- (ii) Let $A = \xi \otimes y$ be such that ξ is a Bessel vector of \mathcal{U} . Then, for any $\eta \in H$, $A\eta = \langle \eta, y \rangle \xi$ is also a Bessel vector of \mathcal{U} and hence A is a Bessel generator multiplier.

For a general unitary system it is very hard to have a good characterization for its Bessel generator multipliers. In this section, we examine several interesting classes of Bessel generator multipliers. We consider a special case of unitary systems when $\mathcal{U} = \mathcal{U}_1 \mathcal{U}_0$, where \mathcal{U}_0 is a countable group (but not necessarily abelian) and $\mathcal{U}_1 \cap U_0 = \{I\}$. These include the group systems, the wavelet systems and the Gabor systems (for the Gabor systems we refer to any standard literature for the definitions). In many cases, we also require \mathcal{U}_1 to be a countable group. For a subset \mathcal{A} of the algebra $B(H)$ of all bounded linear operators on H , we will use $w^*(\mathcal{A})$ to denote the von Neumann algebra generated by \mathcal{A} , and let $\mathcal{A}' = \{S \in B(H) : AS = SA, \forall S \in \mathcal{A}\}$ denote the commutant of \mathcal{A} . Clearly, the set of all Bessel generator multipliers for a unitary system \mathcal{U} is a subalgebra of $B(H)$ with respect to the operator multiplication, and by Example 7.1 every operator in \mathcal{U}' is a Bessel generator multiplier.

Proposition 7.3. *Let $\mathcal{U} = \mathcal{U}_1 \mathcal{U}_0$ be a unitary system such that \mathcal{U}_0 is a countable group and $\mathcal{U}_1 \cap U_0 = \{I\}$. If $A \in w^*(\mathcal{U}_0) \cap \mathcal{U}_0'$, then A is a Bessel generator multiplier for \mathcal{U} .*

Proof. Let ψ be a Bessel vector for \mathcal{U} , let $S_\psi f = \sum_{U_0 \in \mathcal{U}_0} \langle f, U_0 \psi \rangle U_0 \psi$. Then $S_\psi \in \mathcal{U}_0'$. Moreover, since $A \in \mathcal{U}_0'$, we get

$$\begin{aligned} S_{A\psi} f &= \sum_{U_0 \in \mathcal{U}_0} \langle f, U_0 A \psi \rangle U_0 A \psi \\ &= A \sum_{U_0 \in \mathcal{U}_0} \langle A^* f, U_0 \psi \rangle U_0 \psi \\ &= A S_\psi A^* f. \end{aligned}$$

Thus, we have, $S_{A\psi} = A S_\psi A^*$. Let $K = \|A\|^2$. Then $AA^* \leq KI$. By the assumption that $A \in w^*(\mathcal{U}_0)$ and the fact that $S_\psi \in \mathcal{U}_0'$, we obtain $S_\psi^{1/2} A = A S_\psi^{1/2}$. Thus, we have

$$S_{A\psi} = A S_\psi A^* = S_\psi^{1/2} A A^* S_\psi^{1/2} \leq K S_\psi,$$

and so $U S_{A\psi} U^* \leq K U S_\psi U^*$ for any bounded operator U . This implies that

$$\begin{aligned} \sum_{U \in \mathcal{U}} (U A \psi) \otimes (U A \psi) &= \sum_{U_1 \in \mathcal{U}_1, U_0 \in \mathcal{U}_0} (U_1 U_0 A \psi) \otimes (U_1 U_0 A \psi) \\ &= \sum_{U_1 \in \mathcal{U}_1} U_1 S_{A\psi} U_1^* \end{aligned}$$

$$\begin{aligned}
&\leq K \sum_{U_1 \in \mathcal{U}_1} U_1 S_\psi U_1^* \\
&= K \sum_{U_1 \in \mathcal{U}_1, U_0 \in \mathcal{U}_0} (U_1 U_0 \psi) \otimes (U_1 U_0 \psi) \\
&= K \sum_{U \in \mathcal{U}} (U \psi) \otimes (U \psi).
\end{aligned}$$

Since ψ is Bessel, we have that $\sum_{U \in \mathcal{U}} (U \psi) \otimes (U \psi)$ is weakly convergent to a positive bounded linear operator on H , and so $\sum_{U \in \mathcal{U}} (U A \psi) \otimes (U A \psi)$ is also weakly convergent to a positive bounded linear operator. Therefore, $A \psi$ is a Bessel vector for \mathcal{U} , which implies that A is a Bessel vector multiplier, as claimed. \square

Corollary 7.1. *Let $\mathcal{U} = \mathcal{U}_1 \mathcal{U}_0$ be a unitary system such that \mathcal{U}_0 is a group and $\mathcal{U}_1 \cap U_0 = \{I\}$. Assume that \mathcal{U}_0 is abelian. Then every operator in $w^*(\mathcal{U}_0)$ is a Bessel generator multiplier for \mathcal{U} .*

Proposition 7.4. *Let $\mathcal{U} = \mathcal{U}_1 \mathcal{U}_0$ be a unitary system such that \mathcal{U}_0 is a group and $\mathcal{U}_1 \cap U_0 = \{I\}$. Assume that $A \in \mathcal{U}_0'$ is invertible and that $U_1^* A^{-1} U_1 A \in w^*(\mathcal{U}_0)$ for every $U_1 \in \mathcal{U}_1$, then A is a Bessel generator multiplier for \mathcal{U} .*

Proof. Let ψ be a Bessel vector for \mathcal{U} . Fix $U_1 \in \mathcal{U}_1$, let $B = U_1^* A^{-1} U_1 A$ and let S_ψ be defined as in the proof of Proposition 7.3. Since $S_\psi \in \mathcal{U}_0'$, we have $BS^{\frac{1}{2}} = S^{\frac{1}{2}}B$. Thus we get

$$BS_\psi B^* = S^{\frac{1}{2}}(BB^*)S^{\frac{1}{2}} \leq \|B\|^2 \cdot S_\psi \leq K \cdot S_\psi,$$

where $K = (\|A^{-1}\| \cdot \|A\|)^2$. On the other hand we also have

$$BS_\psi B^* = (U_1^* A^{-1} U_1 A) S_\psi (A^* U_1^* (A^{-1})^* U_1) = U_1^* A^{-1} U_1 S_{A\psi} U_1^* (A^{-1})^* U_1$$

Therefore, we get

$$U_1^* A^{-1} U_1 S_{A\psi} U_1^* (A^{-1})^* U_1 \leq K \cdot S_\psi,$$

which implies that

$$U_1 S_{A\psi} U_1^* \leq KA[U_1 S_\psi U_1^*]A^*,$$

holds for every $U_1 \in \mathcal{U}_1$. Since $\sum_{U_1 \in \mathcal{U}_1} U_1 S_\psi U_1^*$ is weakly convergent to a positive bounded linear operator, we have that $\sum_{U_1 \in \mathcal{U}_1} U_1 S_{A\psi} U_1^*$ is also weakly convergent to a positive bounded linear operator. Hence, $A\psi$ is a Bessel vector for \mathcal{U} . \square

7.4.2 The Bessel Generator Multiplier Algebra

Let $\mathcal{M}_{\mathcal{U}}$ be the set of all Bessel generator multipliers for a unitary system \mathcal{U} . Then clearly, $\mathcal{M}_{\mathcal{U}}$ is a subalgebra of $B(H)$, and we will call it the *Bessel generator multiplier algebra* for \mathcal{U} . In general, $\mathcal{M}_{\mathcal{U}}$ is not a closed subalgebra of $B(H)$,

and in fact in many cases it is usually a very large dense subalgebra of $B(H)$ (see Proposition 7.5). Here, we will introduce a natural quantity associated with the Bessel bounds for each operator in $\mathcal{M}_{\mathcal{U}}$. For any Bessel vector η of \mathcal{U} , we will use K_{η} to denote the optimal (minimal) Bessel bound of η . For $A \in \mathcal{M}_{\mathcal{U}}$, we define the *Bessel norm* of A by

$$\|A\|_b = \sup\{\sqrt{K_{A\eta}} : K_{\eta} \leq 1, \eta \in \mathcal{B}_{\mathcal{U}}\}$$

We will show that the Bessel norm is indeed a norm for the Bessel multiplier algebra, and we also compare the two norms for special classes of Bessel generator multipliers.

We already know that $K_{\xi} = \|\Theta_{\xi}\|^2 = \|x\|_o^2$. Therefore, we have

$$\|A\|_b = \sup\{\|A\eta\|_o : \eta \in \mathcal{B}_{\mathcal{U}}, \|\eta\|_o = 1\}.$$

Thus, the Bessel norm $\|\cdot\|_b$ would indeed be a norm if we could prove that $A : (\mathcal{B}_{\mathcal{U}}, \|\cdot\|_o) \rightarrow (\mathcal{B}_{\mathcal{U}}, \|\cdot\|_o)$ is a bounded linear operator. Indeed this can be proved with the closed graph theorem.

Theorem 7.3. *The Bessel norm $\|\cdot\|_b$ is a norm on $\mathcal{M}_{\mathcal{U}}$. Moreover, $(\mathcal{M}_{\mathcal{U}}, \|\cdot\|_b) \cap \{A \in B(H) : \|A\| \leq 1\}$ is a closed subset of $B(\mathcal{B}_{\mathcal{U}}, \|\cdot\|_o)$.*

Proof. Let A be a Bessel generator multiplier. Then it is a linear operator on $\mathcal{B}_{\mathcal{U}}$. Since $\|A\|_b = \sup\{\|A\eta\|_o : \eta \in \mathcal{B}_{\mathcal{U}}, \|\eta\|_o = 1\}$, we only need to show that A is bounded with the norm $\|\cdot\|_o$ equipped on $\mathcal{B}_{\mathcal{U}}$.

Now let $x_n, x, y \in \mathcal{B}_{\mathcal{U}}$ be such that $\|x_n - x\|_o \rightarrow 0$ and $\|Ax_n - y\|_o \rightarrow 0$. We need to show that $y = Ax$. In fact, from the proof of Lemma 7.2, we must have $\|x_n - x\| \rightarrow 0$ and $\|Ax_n - y\| \rightarrow 0$. Thus, we have $y = Ax$ since A is a bounded linear operator on H . Therefore, by the closed graph theorem we get that A is bounded with the norm $\|\cdot\|_o$ equipped on $\mathcal{B}_{\mathcal{U}}$.

For the moreover part, suppose that $A_n \in \mathcal{M}_{\mathcal{U}}$ and $A \in B(\mathcal{B}_{\mathcal{U}}, \|\cdot\|_o)$ are such that $\|A_n - A\|_b \rightarrow 0$ with $\|A_n\| \leq 1$. We will prove that A is a restriction of a bounded linear operator on H , and hence $A \in \mathcal{M}_{\mathcal{U}}$. It suffices to show that $A \in B(\mathcal{B}_{\mathcal{U}}, \|\cdot\|)$. In fact, from $\|A_n - A\|_b \rightarrow 0$, we get

$$\|A_n x - Ax\| \leq \|A_n - A\|_b \cdot \|x\|_o \rightarrow 0$$

for every $x \in \mathcal{B}_{\mathcal{U}}$. This implies that

$$\|Ax\| \leq \sup\{\|A_n x\| : n \in \mathbb{N}\} \leq \sup\{\|A_n\| \cdot \|x\| : n \in \mathbb{N}\} \leq \|x\|$$

for every $x \in \mathcal{B}_{\mathcal{U}}$. So A can be extended to a bounded linear operator, say \tilde{A} , on H with $\|\tilde{A}\| \leq 1$. Thus $(\mathcal{M}_{\mathcal{U}}, \|\cdot\|_b) \cap \{A \in B(H) : \|A\| \leq 1\}$ is a closed subset of $B(\mathcal{B}_{\mathcal{U}}, \|\cdot\|_o)$. \square

In general, $(\mathcal{M}_{\mathcal{U}}, \|\cdot\|_b)$ is not complete (see Proposition 7.8). For operator norm closure we have the following:

Proposition 7.5. *Assume that $\mathcal{B}_{\mathcal{U}}$ is dense in H . Then the closure of $\mathcal{M}_{\mathcal{U}}$ in the usual operator norm contains all the compact operators, and hence, $\mathcal{M}_{\mathcal{U}}$ is dense in $B(H)$ in the strong operator topology. Therefore, if $\mathcal{B}_{\mathcal{U}} \neq H$, then $\mathcal{M}_{\mathcal{U}}$ is not closed in the operator norm on $B(H)$.*

Proof. Given $x, y \in H$. Since $\mathcal{B}_{\mathcal{U}}$ is dense in H , there exists a sequence $\{\eta_n\}$ from $\mathcal{B}_{\mathcal{U}}$ such that $\|\eta_n - x\| \rightarrow 0$. Let $A_n = \eta_n \otimes y$ and $A = x \otimes y$. Then $\|A_n - A\| \rightarrow 0$ and $A_n \in \mathcal{M}_{\mathcal{U}}$. Thus, the operator norm closure of $\mathcal{M}_{\mathcal{U}}$ contains all the compact operators. Hence, $\mathcal{M}_{\mathcal{U}}$ is dense in $B(H)$ in the strong operator topology. In the case that $\mathcal{B}_{\mathcal{U}} \neq H$, let $x \notin \mathcal{B}_{\mathcal{U}}$ and $y \neq 0$. Then from the above argument $A = x \otimes y$ is in the operator norm closure of $\mathcal{M}_{\mathcal{U}}$, but clearly $A \notin \mathcal{M}_{\mathcal{U}}$. Thus $\mathcal{M}_{\mathcal{U}}$ is not closed in the operator norm on $B(H)$. \square

Proposition 7.6. (a) *If $A \in \mathcal{U}'$, then $\|A\|_b \leq \|A\|$. Moreover, if in addition, we assume that \mathcal{U} admits a Parseval frame vector, then we have $\|A\|_b = \|A\|$.*
 (b) *If $A \in w^*(\mathcal{U}_0) \cap \mathcal{U}'_0$, we have $\|A\|_b \leq \|A\|$.*
 (c) *If A is a multiplier satisfying the condition in Proposition 7.4, then $\|A\|_b \leq \|A^{-1}\| \cdot \|A\|^2$.*

Proof. (a) Let $x \in \mathcal{B}_{\mathcal{U}}$. Then

$$\|Ax\|_o = \|\Theta_{Ax}\| = \|\Theta_x A^*\| \leq \|A^*\| \cdot \|\Theta_x\| = \|A\| \cdot \|x\|_o.$$

Thus, $\|A\|_b \leq \|A\|$. For the moreover part, let ξ be a Parseval frame vector. Then Θ_{ξ} is an isometry and $\|\xi\|_0 = 1$. Thus

$$\|A\|_b \geq \|A\xi\|_o = \|\Theta_{A\xi}\| = \|\Theta_{\xi} A^*\| = \|A^*\| = \|A\|,$$

and so we have the equality.

(b) Follows from the proof of Proposition 7.3.

(c) Follows from the proof of Proposition 7.4. \square

7.4.3 The Wavelet Unitary System

We show that for the wavelet unitary system, any of the following three cases could happen: (i) $\|A\|_b = \|A\|$; (ii) $\|A\|_b > \|A\|$; and (iii) $\|A\|_b < \|A\|$. Recall that the wavelet unitary system is the unitary system $\mathcal{U}_{D,T} := \{D^n T^\ell : n, \ell \in \mathbb{Z}\}$, where D is the dilation operator and T is the translation unitary operator defined by

$$(Df)(x) = \sqrt{2}f(2x), \quad (Tf)(x) = f(x-1), \quad f \in L^2(\mathbb{R}).$$

The Fourier transform we will use here is the normalized one:

$$(\mathcal{F}f)(\xi) = \hat{f}(\xi) = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} f(x) e^{-ix\xi} dx$$

for all $f \in L^1(\mathbb{R})$ and then extend it unitarily to $L^2(\mathbb{R})$. With this definition we have $\hat{D} = \mathcal{F}D\mathcal{F}^{-1} = D^{-1}$ and $\hat{T} = \mathcal{F}T\mathcal{F}^{-1} = M_{e^{-i\xi}}$, where $M_{e^{-i\xi}}$ is the multiplication operator by $e^{-i\xi}$. It is known (cf. [8]) that the commutant of $\{\hat{D}, \hat{T}\}$ is $\{M_h : h \in L^\infty(\mathbb{R}), h(2\xi) = h(\xi), a.e., \xi \in \mathbb{R}\}$. A Parseval frame set is a measurable subset F of \mathbb{R} such that $\{2^n F : n \in \mathbb{Z}\}$ is a partition of \mathbb{R} (modulo measure zero sets) and $\{F + 2n\pi : n \in \mathbb{Z}\}$ is disjoint (modulo measure zero sets). Equivalently, F is a Parseval frame set if and only if $\{\hat{D}^n \hat{T}^\ell \left(\frac{1}{\sqrt{2\pi}} \chi_F \right) : n, \ell \in \mathbb{Z}\}$ is a Parseval frame set for $L^2(\mathbb{R})$ (c.f. [20]). By Proposition 7.6 (a) we have that $\|A\|_b = \|A\|$ if $A \in \{D, T\}'$. So we only need to show that the other two cases could happen for the wavelet unitary system.

Example 7.2. There exists a Bessel generator multiplier for the wavelet unitary system such that $\|A\|_b > \|A\|$.

Proof. Let $A = \psi \otimes \varphi$ be such that ψ is a Bessel wavelet and $\psi \in L^2(\mathbb{R})$. Then, for every $f \in L^2(\mathbb{R})$, $Af = \langle f, \varphi \rangle \psi$ is also a Bessel wavelet and hence A is a Bessel wavelet multiplier with Bessel bound $K_{Af} = |\langle f, \varphi \rangle|^2 K_\psi$.

Now let $\hat{\psi}(\xi) = \frac{1}{\sqrt{2\pi}} \chi_{[-\pi, -\pi/2) \cup [\pi/2, \pi]}$ and $\hat{\varphi}(\xi) = \frac{1}{\sqrt{2\pi}} \chi_{[-2\pi, -\pi) \cup [\pi, 2\pi]}$. Then ψ is a Bessel vector such that $\|\psi\| < \|\psi\|_o = 1$ and φ is an orthonormal wavelet such that $\|\varphi\| = \|\varphi\|_o = 1$. Hence, we get

$$\|A\|_b^2 \geq \|\Theta_{A\varphi}\|^2 = |\langle \varphi, \varphi \rangle|^2 K_\psi = K_\psi = \|\psi\|_o^2 > \|\psi\|^2 = \|A\|^2$$

since $\|A\| = \|\varphi\| \|\psi\| = \|\psi\|$, and so we have $\|A\|_b > \|A\|$. \square

Example 7.3. There exists a Bessel generator multiplier for the wavelet unitary system such that $\|A\|_b < \|A\|$.

Proof. Let $E = [\frac{\pi}{4}, \frac{\pi}{2}]$, and define A by $(\hat{A}f)(\xi) = \chi_{2E}(\xi) f(\frac{\xi}{2})$. We prove that $\|A\| = \sqrt{2}$ and $\|A\|_b = 1$, and consequently we get $\|A\|_b < \|A\|$.

Assume that $f \in L^2(\mathbb{R})$ is such that $\text{supp}(\hat{f}) \subseteq E$. Then

$$\|Af\|^2 = \int_{2E} \left| \chi_{2E}(\xi) \hat{f}\left(\frac{\xi}{2}\right) \right|^2 d\xi = 2 \int_E |\hat{f}(t)|^2 dt = 2\|\hat{f}\|^2 = 2\|f\|^2.$$

Thus, $\|A\| = \sqrt{2}$.

Next, we show that $\|A\|_b = 1$. We will need two facts.

Fact 1: Let g be a Bessel wavelet for the wavelet system such that $\text{supp}(\hat{g})$ is contained in a Parseval frame set F . Then we claim that $\|g\|_o = \sqrt{2\pi} \|\hat{g}\|_\infty$. In fact, since

F is Parseval frame set, we have that $\{2^n F : n \in \mathbb{Z}\}$ is a disjoint partition of \mathbb{R} . So we can extend the restriction of \hat{g} on F to a 2-dilation function \hat{h} (i.e. $\hat{h}(2\xi) = \hat{h}(\xi)$, a.e. $\xi \in \mathbb{R}$). Then $\|\hat{g}\|_\infty = \|\hat{h}\|_\infty$, and the multiplication operator $M_{\hat{h}}$ commutes with both operators \hat{D} and \hat{T} . Let $\hat{\psi} = \frac{1}{\sqrt{2\pi}}\chi_F$. Then $\{\psi_{k,j} : k, j \in \mathbb{Z}\}$ is a Parseval frame for $L^2(\mathbb{R})$, and so $\sum_{k,j \in \mathbb{Z}} \hat{D}^k \hat{T}^j \hat{\psi} \otimes \hat{D}^k \hat{T}^j \hat{\psi} = I$. Note that $\hat{g}\hat{\psi} = \hat{h}\hat{\psi}$. So we get

$$\begin{aligned} S_{\hat{g}} &= \sum_{k,j \in \mathbb{Z}} \hat{D}^k \hat{T}^j \hat{g} \otimes \hat{D}^k \hat{T}^j \hat{g} \\ &= 2\pi \sum_{k,j \in \mathbb{Z}} \hat{D}^k \hat{T}^j M_{\hat{g}} \hat{\psi} \otimes \hat{D}^k \hat{T}^j M_{\hat{g}} \hat{\psi} \\ &= 2\pi \sum_{k,j \in \mathbb{Z}} \hat{D}^k \hat{T}^j M_{\hat{h}} \hat{\psi} \otimes \hat{D}^k \hat{T}^j M_{\hat{h}} \hat{\psi} \\ &= 2\pi M_{\hat{h}} \left(\sum_{k,j \in \mathbb{Z}} \hat{D}^k \hat{T}^j \hat{\psi} \otimes \hat{D}^k \hat{T}^j \hat{\psi} \right) M_{\hat{h}}^* \\ &= 2\pi M_{|\hat{h}|^2}. \end{aligned}$$

So we have $\|g\|_0 = \sqrt{\|S_{\hat{g}}\|} = \sqrt{2\pi} \|\hat{h}\|_\infty = \sqrt{2\pi} \|\hat{g}\|_\infty$ since $\|\hat{g}\|_\infty = \|\hat{h}\|_\infty$.

Fact 2: If g is a Bessel wavelet, then $\sqrt{2\pi} \|\hat{g}\|_\infty \leq \|g\|_0$. It is a well known fact [24] that for any compactly supported function \hat{f} , we have

$$\int_{\mathbb{R}} |\hat{f}(\xi)|^2 \sum_{k \in \mathbb{Z}} |\hat{g}(2^k \xi)|^2 d\xi \leq \frac{1}{2\pi} \sum_{k,j \in \mathbb{Z}} |\langle \hat{f}, \hat{D}^k \hat{T}^j \hat{g} \rangle|^2$$

Therefore, we get

$$\int_{\mathbb{R}} |\hat{f}(\xi)|^2 \sum_{k \in \mathbb{Z}} |\hat{g}(2^k \xi)|^2 d\xi \leq \frac{1}{2\pi} \|g\|_0^2 \cdot \|\hat{f}\|^2.$$

Thus, $\sum_{k \in \mathbb{Z}} |\hat{g}(2^k \xi)|^2 \leq \frac{1}{2\pi} \|g\|_0^2$ for a.e. $\xi \in \mathbb{R}$, which implies that $\sqrt{2\pi} \|\hat{g}\|_\infty \leq \|g\|_0$.

Now let g be any Bessel wavelet for the wavelet system. Then $\text{supp}(\hat{A}\hat{g}) \subseteq 2E$. Since $2E$ is a subset of the Shannon wavelet set $[-2\pi, -\pi) \cup [\pi, 2\pi)$, we get from Fact 1 that $\|Ag\|_0 = \sqrt{2\pi} \|\hat{A}\hat{g}\|_\infty \leq \sqrt{2\pi} \|\hat{g}\|_\infty$. However, by Fact 2, we also have $\sqrt{2\pi} \|\hat{g}\|_\infty \leq \|g\|_0$. Therefore, we obtain that $\|Ag\|_0 \leq \|g\|_0$ and hence $\|A\|_b \leq 1$. On the other hand, if g be a Bessel wavelet for the wavelet system such that $\hat{g}(\xi) = \chi_E(\xi)$. Then we get $\|Ag\|_0 = \sqrt{2\pi} \|\hat{g}\|_\infty = \|g\|_0$. Thus, $\|A\|_b = 1$, as claimed. \square

A Bessel wavelet multiplier is called a *Fourier–Bessel wavelet multiplier* if it is of the form $\hat{A} = M_h$ for some $h \in L^\infty(\mathbb{R})$. Note that none of the Bessel wavelet

multipliers in the last two example is a Fourier–Bessel wavelet multiplier. For general Fourier–Bessel wavelet multipliers we have the following:

Example 7.4. Assume that A , with $\hat{A} = M_h$, is a Fourier–Bessel wavelet multiplier for the wavelet unitary system. Then $\|A\| \leq \|A\|_b$.

Proof. For any $\varepsilon > 0$, we can find measurable subset E with positive measure such that $|h(\xi)| \geq \|h\|_\infty - \varepsilon$ on E , and E is both 2-dilation disjoint and 2π -translation disjoint. Let $\hat{\psi} = \frac{1}{\sqrt{2\pi}}\chi_E$. Then ψ is a Bessel vector with Bessel constant 1, i.e. $\|\psi\|_o = 1$. Let $\varphi = A\psi$ (i.e. $\hat{\varphi} = h\hat{\psi}$). We calculate the Bessel bound of φ . Note that $\langle \hat{\varphi}, (h\hat{\psi})_{j,k} \rangle = 0$ if $j \neq 0$. So we have

$$\begin{aligned} \sum_{j,k \in \mathbb{Z}} |\langle \varphi, (h\psi)_{j,k} \rangle|^2 &= \sum_{k \in \mathbb{Z}} |\langle \varphi, (h\psi)_{0,k} \rangle|^2 \\ &= \|\varphi\|^2 = \frac{1}{2\pi} \int_E |h(\xi)|^2 d\xi \\ &\geq (\|h\|_\infty - \varepsilon) \|\varphi\|^2. \end{aligned}$$

This φ has Bessel bound no less than $\|h\|_\infty - \varepsilon$, which implies that $\|A\|_b \geq \|h\|_\infty - \varepsilon$ for all small $\varepsilon > 0$. This implies that

$$\|A\|_b \geq \|h\|_\infty = \|A\|,$$

as claimed. □

We have the following open problem:

Problem A. Is it true that $\|A\|_b = \|A\|$ for all Fourier–Bessel multipliers A ?

We conjecture that the answer is affirmative and point out the following partial result.

Corollary 7.2. Assume that h is a Fourier–Bessel wavelet multiplier for the wavelet unitary system such that \hat{h} is either 2π -translation or 2-dilation periodic. Then for $A = \hat{M}_h$ we have $\|A\| = \|A\|_b$.

Proof. From the proof of Proposition 2.1 in [22] we get that $\|A\|_b \leq \|\hat{h}\|_\infty = \|A\|$. Thus, the conclusion follows from Example 7.4. □

7.4.4 The Group Unitary System Case

In this section, we discuss the case when the unitary system is the image of a projective unitary representation of a countable group. Recall (cf. [38]) that a *projective unitary representation* π for a countable group G is a mapping $g \mapsto \pi(g)$ from G into the group $U(H)$ of all the unitary operators on a separable Hilbert space H

such that $\pi(g)\pi(h) = \mu(g, h)\pi(gh)$ for all $g, h \in G$, where $\mu(g, h)$ is a scalar-valued function on $G \times G$ taking values in the circle group \mathbb{T} . This function $\mu(g, h)$ is then called a *multiplier* of π . In this case, we also say that π is a μ -projective unitary representation. It is clear from the definition that we have

- (i) $\mu(g_1, g_2 g_3) \mu(g_2, g_3) = \mu(g_1 g_2, g_3) \mu(g_1, g_2)$ for all $g_1, g_2, g_3 \in G$,
 - (ii) $\mu(g, e) = \mu(e, g) = 1$ for all $g \in G$, where e denotes the group unit of G .
- Any function $\mu : G \times G \rightarrow \mathbb{T}$ satisfying (i) – (ii) above will be called a *multiplier* for G . It follows from (i) and (ii) that we also have
- (iii) $\mu(g, g^{-1}) = \mu(g^{-1}, g)$ holds for all $g \in G$.

Examples of projective unitary representations include unitary group representations and Gabor representations in time-frequency analysis. In what follows, we will consider the unitary system $\pi(G)$, and use \mathcal{B}_π to denote the set of all Bessel vectors for $\pi(G)$, where π is a projective unitary representation of G .

Lemma 7.3. [9] *Let π be a projective representation of a countable group G on a Hilbert space H and x is a Bessel vector for π . Then there exists $\xi \in M := \overline{\text{span}}\{\pi(g)x : g \in G\}$ such that $\{\pi(g)\xi\}_{g \in G}$ is a Parseval frame for M .*

Lemma 7.4. [18] *Assume that π is a projective representation of a countable group G on a Hilbert space H and ξ is a complete Parseval frame vector for π . Let $\eta \in H$. Then we have*

- (i) η is a Bessel vector for π if and only if there exists an operator $T \in w^*(\pi(G))$ such that $\eta = T\xi$.
- (ii) η is a complete frame vector for π if and only if there exists an invertible operator $T \in w^*(\pi(G))$ such that $\eta = T\xi$.
- (iii) η is a complete Parseval frame vector for π if and only if there exists a unitary operator $T \in w^*(\pi(G))$ such that $\eta = T\xi$.

Proposition 7.7. *Let π be a projective unitary representation of G . If $A \in w^*(\pi(G))$, then A is a Bessel generator multiplier.*

Proof. Let ξ be a Bessel vector for π . Then, by Lemma 7.3, there exists $\eta \in M := \overline{\text{span}}\{\pi(g)\xi : g \in G\}$ such that $\{\pi(g)\eta : g \in G\}$ is a Parseval frame for M . Thus, it follows from Lemma 7.4 that a vector $x \in M$ is a Bessel vector if and only if there exists $B \in w^*(\pi(G))|_M$ such that $x = B\eta$. Let $\xi = T\eta$ with $T \in w^*(\pi(G))|_M$. Then $A\xi = AT\eta$ with $AT \in w^*(\pi(G))|_M$ is also Bessel. Therefore, A is a Bessel generator multiplier. \square

Theorem 7.4. *Let π be a projective unitary representation of G such that \mathcal{B}_π is dense in H . Then*

- (i) $\|A\|_b = \|A\|$ if $A \in \pi(G)'$
- (ii) $\|A\|_b \leq \|A\|$ if $A \in w^*(\pi(G))$.

We remark that part (i) of this theorem is not the consequence of Proposition 7.6 since here we do not assume π admits a Parseval frame vector for H .

In order to prove Theorem 7.4, we need the following lemma:

Lemma 7.5. *Let π be a projective unitary representation of G such that its has a complete wandering vector ψ . Then $\|A\|_b = \|A\|$ when $A \in w^*(\pi(G))$.*

Proof. Let $\sigma : w^*(\pi(G)) \rightarrow \pi(G)'$ be the conjugate linear anti-automorphism such that $T\psi = \sigma(T)\psi$ for all $T \in w^*(\pi(G))$ (cf. [28]). Let $A \in w^*(\pi(G))$. Then

$$\|A\|_b \geq \|\Theta_{A\psi}\| = \|\Theta_{\sigma(A)\psi}\| = \|\Theta_{\psi\sigma(A)^*}\| = \|\sigma(A)^*\| = \|A\|,$$

where we use the isometry property of Θ_{ψ} .

On the other hand, for any Bessel vector $x \in \mathcal{B}_{\pi}$, there exists $B \in \pi(G)'$ such that $x = B\psi$. Thus, we have

$$\begin{aligned} \|Ax\|_b &= \|\Theta_{Ax}\| = \|\Theta_{AB\psi}\| = \|\Theta_{BA\psi}\| = \|\Theta_{B\sigma(A)\psi}\| \\ &= \|\Theta_{\psi\sigma(A)^*B^*}\| \leq \|\sigma(A)^*\| \cdot \|\Theta_{\psi}\| \cdot \|B^*\| \\ &= \|A\| \cdot \|\Theta_{\psi}\| \cdot \|B\| = \|A\| \cdot \|B\|. \end{aligned}$$

Note that $\|x\|_o = \|\Theta_x\| \cdot \|\Theta_{B\psi}\| = \|B^*\| = \|B\|$. Thus, we get

$$\|Ax\|_b \leq \|A\| \cdot \|x\|_o$$

for all $x \in \mathcal{B}_{\pi}$. This implies that $\|A\|_b \leq \|A\|$, as claimed. \square

Lemma 7.6. *Let π be a μ -projective unitary representation of G on a Hilbert space H such that its has a complete frame vector ψ . Then there exists a μ -projective unitary representation σ of G on a Hilbert space K such that the μ -projective unitary representation of $\pi \oplus \sigma$ on $H \oplus K$ has a complete wandering vector.*

Corollary 7.3. *Let π be a projective unitary representation of G such that its has a complete frame vector. Then $\|A\|_b \leq \|A\|$ when $A \in w^*(\pi(G))$.*

Proof. By Lemma 7.6, we can dilate π to $\pi \oplus \sigma$ on $H \oplus K$ for some Hilbert space K such that $\pi \oplus \sigma$ has complete wandering vector. Let $\tilde{A} \in w^*(\pi(g) \oplus \sigma(g) : g \in G)$ such that $\tilde{A}|_H = A$ and $\|\tilde{A}\| = \|A\|$. By Lemma 7.5, we have $\|\tilde{A}\|_b = \|\tilde{A}\|$. Thus $\|A\|_b \leq \|\tilde{A}\|_b = \|\tilde{A}\| = \|A\|$. \square

Proof of Theorem 7.4. (i) For any $\xi \in \mathcal{B}_{\pi}$, we have

$$\sum_{g \in G} |\langle x, \pi(g)Ag \rangle|^2 = \sum_{g \in G} |\langle A^*x, \pi(g)g \rangle|^2 \leq \|A^*\|^2 K_g \|x\|, \quad x \in H$$

where K_g is Bessel bound of g . Thus, we have $\|A\|_b \leq \|A\|$. So we only need to show that $\|A\|_b \geq \|A\|$. Since \mathcal{B}_{π} is dense in H , for each $\varepsilon > 0$ we can find $x \in \mathcal{B}_{\pi}$

with $\|x\| = 1$ and $\|A^*x\| \geq \|A\| - \varepsilon$. Let A^*x , and $H_x = [\pi(G)A^*x]$. Then, by Lemma 7.3, there exists a $\xi \in H_x$ such that $\{\pi(g)\xi\}$ is a Parseval frame for H_y . So $\|\xi\|_o = 1$, and

$$\|A\|_b \geq \|\Theta_{A\xi}\| = \|\Theta_\xi A^*\| \geq \|\Theta_\xi A^*x\| = \|A^*x\| \geq \|A\| - \varepsilon,$$

where we use that Θ_ξ is an isometry when restricted to H_y and $A^*x \in H_y$. Thus, $\|A\|_b \geq \|A\|$ by letting $\varepsilon \rightarrow 0$.

(ii) Let $x \in \mathcal{B}_\pi$ such that $\|x\|_o \leq 1$. Let $M = [\pi(G)x]$ and consider $\pi|_M$. Then, by Corollary 7.3, we get $\|\bar{A}\|_b \leq \|\bar{A}\|$. This implies that

$$\|Ax\|_o = \|\bar{A}x\|_o \leq \|\bar{A}\|_b \leq \|\bar{A}\| \leq \|A\|.$$

Taking the sup over $\|x\|_o \leq 1$, we get

$$\|A\|_b \leq \|A\|$$

as claimed. \square

Problem B. Let π be a projective unitary representation of G . Do we have $\|A\|_b = \|A\|$ for $A \in w^*(\pi(G))$? What about other Bessel generator multipliers?

Let A be a Bessel generator multiplier for a general unitary system \mathcal{U} . Then, by Theorem 7.3, we have that A is also a bounded linear mapping Φ_A from the operator space $(\mathcal{B}_\mathcal{U}, \|\cdot\|_o)$ into itself. One of the most important concept in the theory of operator spaces is the completely bounded maps. Recall that a linear mapping Φ from an operator space S to an operator space V is called completely bounded if $\sup\{\|\Phi_n\| : n \geq 1\} < \infty$, where Φ_n is defined by

$$\Phi_n([a_{ij}]) = [\Phi(a_{ij})]$$

on the operator space $M_n \otimes S$. A natural question about the Bessel vector space is to examine the completely boundedness of Φ_A . We ask the following:

Problem C. Is Φ_A always completely bounded?

We examine a special case of this question. Assume that \mathcal{U} admits a complete wandering vector, say ψ . Then, by Lemma 7.2, we can identify $(\mathcal{B}_\mathcal{U}, \|\cdot\|_o)$ with the local commutant operator space $\mathcal{C}_\psi(\mathcal{U})$. Therefore, if A is a Bessel generator multiplier for \mathcal{U} , then induced mapping Φ_A is defined by $\Phi_A(S) = T$ where T is the unique operator in $\mathcal{C}_\psi(\mathcal{U})$ such that $T\psi = AS\psi$. However, in general we don't have $T = AS$ since AS may not be in $\mathcal{C}_\psi(\mathcal{U})$. In the case that $A \in \mathcal{U}'$, then $AS \in \mathcal{C}_\psi(\mathcal{U})$ for every $S \in \mathcal{C}_\psi(\mathcal{U})$, and therefore $\Phi_A(S) = AS$, and so Φ is completely bounded. However, it remains open whether all the induced mapping Φ_A is completely bounded even when \mathcal{U} admits a complete wandering vector.

The second natural question is to ask that if every bounded linear mapping from the operator space $(\mathcal{B}_\mathcal{U}, \|\cdot\|_o)$ to itself is induced by a Bessel generator multiplier from $B(H)$. The following result answers this question negatively.

Proposition 7.8. *Let \mathcal{U} be the wavelet unitary system. Then there exists a bounded linear mapping Φ from Bessel generator multiplier the operator space $(\mathcal{B}_{\mathcal{U}}, \|\cdot\|_o)$ to itself such that is the limit of a Bessel generator multipliers in the Bessel norm. But Φ is unbounded with respect to the Hilbert space norm on $\mathcal{B}_{\mathcal{U}}$. In particular, the Bessel generator multiplier algebra $(\mathcal{M}_{\mathcal{U}}, \|\cdot\|_b)$ is not complete.*

Proof. Let $E_n = [\pi/2^{n+1}, \pi/2^n)$ for $n \geq 1$. Define A_n by $\hat{A}_n f = \chi_{2^n E_n}(\xi) f(\frac{\xi}{2^n})$. Then, as in the proof of Example 7.3, we get that $\|A_n\| = 2^{n/2}$, and $\|A_n\|_b = 1$ for all $n \geq 1$. Now let $B_n = 2^{-n/4} A_n$. Then $\sum_{n=1}^{\infty} B_n$ converges in the $\|\cdot\|_b$ norm on $(\mathcal{B}_{\mathcal{U}}, \|\cdot\|_b)$. Let $\Phi = \sum_{n=1}^{\infty} \hat{B}_n$.

We claim that there is no Bessel generator multiplier A for \mathcal{U} such that $\Phi_A = \Phi$, i.e., $\Phi \neq A$ on $\mathcal{B}_{\mathcal{U}}$ for all Bessel generator multiplier A for \mathcal{U} . In order to prove this, it suffices to show that Φ is an unbounded linear operator with respect to the Hilbert space norm on $(\mathcal{B}_{\mathcal{U}})$. In fact, let $f_n = \chi_{E_n}$. Since $\hat{B}_m f_n = 0$ if $m \neq n$, we have $\Phi(f_n) = \hat{B}_n f_n$. So

$$\|\Phi(f_n)\|^2 = \|\hat{B}_n f_n\|^2 = 2^{-n/2} \mu(2^n E_n) = 2^{-n/2} \left(\frac{\pi}{2}\right).$$

But $\|f_n\|^2 = \mu(E_n) = \frac{\pi}{2^{n+1}}$. Thus, $\|\Phi(f_n/\|f_n\|)\|^2 = 2^{n/2} \rightarrow \infty$. Hence, Φ is unbounded as claimed. \square

References

1. R. Balan, A study of Weyl-Heisenberg and wavelet frames, Ph. D. Thesis, Princeton University, 1998.
2. M. Bownik, Connectivity and density in the set of framelets *Math. Res. Lett.*, 14 (2007), 285–293.
3. M. Bownik, The closure of the set of tight frame wavelets, *Acta Appl. Math.*, 107 (2009), 195–201.
4. X. Dai and Y. Diao, The path-connectivity of s -elementary tight frame wavelets, *J. Appl. Funct. Anal.*, 2 (2007), no. 4, 309–316.
5. X. Dai, Y. Diao and Q. Gu, Frame wavelets with frame set support in the frequency domain, *Illinois J. Math.*, 48 (2004), no. 2, 539–558.
6. X. Dai, Y. Diao, Q. Gu and D. Han, The S -elementary frame wavelets are path connected, *Proc. Amer. Math. Soc.*, 132 (2004), no. 9, 2567–2575.
7. X. Dai, Y. Diao, Q. Gu and D. Han, Frame wavelets in subspaces of $L^2(\mathbb{R}^d)$, *Proc. Amer. Math. Soc.*, 130 (2002), no. 11, 3259–3267.
8. X. Dai and D. Larson, Wandering vectors for unitary systems and orthogonal wavelets, *Mem. Amer. Math. Soc.*, 134 (1998), no. 640.
9. J. Gabardo and D. Han, Subspace Weyl-Heisenberg frames, *J. Fourier Anal. Appl.*, 7 (2001), 419–433.
10. J-P. Gabardo and D. Han, Aspects of Gabor analysis and operator algebras. Advances in Gabor analysis, 129–152, Appl. Numer. Harmon. Anal., Birkhäuser Boston, Boston, MA, 2003
11. J-P. Gabardo and D. Han, Frame representations for group-like unitary operator systems, *J. Operator Theory*, 49 (2003), 223–244.
12. J-P. Gabardo and D. Han, The uniqueness of the dual of Weyl-Heisenberg subspace frames, *Appl. Comput. Harmon. Anal.*, 17 (2004), no. 2, 226–240.

13. J-P. Gabardo, D. Han and D. Larson, Gabor frames and operator algebras, *Wavelet Applications in Singnal and Image Analysis, Proc. SPIE.*, 4119 (2000), 337–345.
14. G. Garrigós, Connectivity, homotopy degree, and other properties of α -localized wavelets on R , *Publ. Mat.*, 43 (1999), no. 1, 303–340.
15. G. Garrigós, E. Hernández, H. Šikić and F. Soria, Further results on the connectivity of Parseval frame wavelets, *Proc. Amer. Math. Soc.*, 134 (2006), no. 11, 3211–3221.
16. G. Garrigós, E. Hernández, H. Šikić, F. Soria, G. Weiss and E. Wilson, Connectivity in the set of tight frame wavelets (TFW), *Glas. Mat. Ser. III.* 38(58) (2003), no. 1, 75–98.
17. G. Garrigós and D. Speegle, Completeness in the set of wavelets, *Proc. Amer. Math. Soc.*, 128 (2000), no. 4, 1157–1166.
18. D. Han, Approximations for Gabor and wavelet frames, *Trans. Amer. Math. Soc.*, 355 (2003), no. 8, 3329–3342.
19. D. Han, Frame Representations and Parseval Duals with Applications to Gabor Frames, *Trans. Amer. Math. Soc.*, 360(2008), 3307–3326.
20. D. Han and D. Larson, Frames, bases and group parametrizations, *Mem. Amer. Math. Soc.*, 697 (2000).
21. D. Han and D. Larson, Wandering vector multipliers for unitary groups, *Trans. Amer. Math. Soc.*, 353(2001), 3347–3370.
22. D. Han and D. Larson, On the orthogonality of frames and the density and connectivity of wavelet frames, *Acta Appl. Math.*, 107 (2009), 211–222.
23. D. Han, Q. Sun and W. Tang, Topological and geometric properties of refinable functions and MRA affine frames, *Appl. Comput. Harm. Anal.*, to appear.
24. E. Hernández, and G. Weiss, *A First Course on Wavelets*, CRC Press, 1996.
25. E. Hernández, X. Wang and G. Weiss, Smoothing minimally supported frequency wavelets. I, *J. Fourier Anal. Appl.*, 2 (1996), no. 4, 329–340.
26. E. Hernández, X. Wang and G. Weiss, Smoothing minimally supported frequency wavelets. II., *J. Fourier Anal. Appl.*, 3 (1997), no. 1, 23–41.
27. G. Ji and K. Saito, On wandering vector multipliers for unitary groups, *Proc. Amer. Math. Soc.*, 133 (2005), 3263–3269.
28. R. Kadison and J. Ringrose, *Fundamentals of the Theory of Operator Algebras, Vol. I and II*, Academic Press, Inc. 1983 and 1985.
29. D. Larson, Von Neumann algebras and wavelets. Operator algebras and applications (Samos, 1996), 267–312, *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, 495, Kluwer Acad. Publ., Dordrecht, 1997.
30. D. Larson, Frames and wavelets from an operator-theoretic point of view. Operator algebras and operator theory (Shanghai, 1997), 201–218, *Contemp. Math.*, 228, Amer. Math. Soc., Providence, RI, 1998.
31. D. Larson, Unitary systems and wavelet sets. Wavelet analysis and applications, 143–171, *Appl. Numer. Harmon. Anal.*, Birkhäuser Basel, 2007.
32. D. Larson, Unitary systems, wavelet sets, and operator-theoretic interpolation of wavelets and frames, *Gabor and Wavelet Frames*, World Scientific (2007), 166–214.
33. R. Liang, Wavelets, their phases, multipliers and connectivity, Ph. D. Thesis, University of North Carolina-Charlotte, 1998.
34. M. Paluszynski, H. Šikić, G. Weiss and S. Xiao, Tight frame wavelets, their dimension functions, MRA tight frame wavelets and connectivity properties, *Adv. in Comp. Math.*, **18** (2003), 297–327.
35. G. Pisier, Introduction to Operator Space Theory, London Math. Soc. Lecture Notes. Cambridge Univ. Press (2003).
36. D. Speegle, Ph. D. thesis, Teaxs A& M University, 1997.
37. D. Speegle, The s -elementary wavelets are path-connected, *Proc. Amer. Math. Soc.*, Vol 127 (1999), 223–233.
38. V. S. Varadarajan, *Geometry of Quantum Theory*, Second Edition, Springer-Verlag, New York–Berlin, 1985.
39. Wutam Consortium, Basic Properties of Wavelets, *J. Fourier Anal. Appl.*, Vol 4, no. 4–5 (1998), 575–594

Chapter 8

The Zak Transform(s)

Eugenio Hernández, Hrvoje Šikić, Guido L. Weiss, and Edward N. Wilson

Abstract The Zak transform has been used in engineering and applied mathematics for several years and many purposes. In this paper, we show how it can be used to obtain an exceedingly elementary proof of the Plancherel theorem and for developing many results in Harmonic Analysis in particularly direct and simple ways. Many publications state that it was introduced in the middle sixties. It is remarkable that only a small number of mathematicians know this and that many textbooks continue to give much harder and less transparent proofs of these facts. We cite a 1950 paper by I. Gelfand and a book by A. Weil, written in 1940 that indicate that in a general non-compact LCA setting the Fourier transform is an average of Zak transforms (which are really Fourier series expressions). We actually introduce versions of these transforms that show how naturally and simply one obtains these results.

8.1 Introduction

We introduce the operator Z that is often called the *Zak Transform*. Our definition is a bit different from the one that usually appears in the literature. We will discuss this difference and will also give a historical account that the reader may find particularly interesting. In order to do this, however, we need to present our treatment of the operator Z (and \tilde{Z}) which shows that the Fourier transform and its inverse are unitary as an immediate consequence of the basic properties of Fourier series.

The operator Z maps each $f \in L^2(\mathbf{R})$ into the function

$$(Zf)(x, \xi) = \sum_{k \in \mathbf{Z}} f(x+k) e^{-2\pi i k \xi} \equiv \varphi(x, \xi), \quad (8.1)$$

H. Šikić (✉)

Department of Mathematics, University of Zagreb, Bijenička 30, HR-10000,
Zagreb, Croatia
e-mail: hsikic@math.hr

$x, \xi \in \mathbf{R}$. Let us explain the meaning of this equality. Since $f \in L^2(\mathbf{R})$,

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \sum_{k \in \mathbf{Z}} |f(x+k)|^2 dx = \int_{-\infty}^{\infty} |f(x)|^2 dx < \infty.$$

Thus,

$$\sum_{k \in \mathbf{Z}} |f(x+k)|^2 < \infty \quad (8.2)$$

for a.e. $x \in \mathbf{R}$. This means that for a.e. $x \in \mathbf{R}$ the series in (8.1) is the Fourier series of a function in $L^2([-\frac{1}{2}, \frac{1}{2}])$ (considered to be 1-periodic in ξ) we denote by $\varphi(x, \xi)$. Moreover, for a.e. $x \in \mathbf{R}$

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} |\varphi(x, \xi)|^2 d\xi = \sum_{k \in \mathbf{Z}} |f(x+k)|^2.$$

It follows, therefore, that

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} \int_{-\frac{1}{2}}^{\frac{1}{2}} |\varphi(x, \xi)|^2 d\xi dx = \int_{-\frac{1}{2}}^{\frac{1}{2}} \sum_{k \in \mathbf{Z}} |f(x+k)|^2 dx = \|f\|_2^2. \quad (8.3)$$

Z , therefore, maps $L^2(\mathbf{R})$ isometrically into a space of functions

$$\varphi \in L^2(\mathbf{T}^2) = L^2\left(\left[-\frac{1}{2}, \frac{1}{2}\right] \times \left[-\frac{1}{2}, \frac{1}{2}\right]\right).$$

Let us examine the space M of these images $\varphi(x, \xi) = (Zf)(x, \xi)$ of $L^2(\mathbf{R})$ under the transformation Z . We have seen that these images are functions of two real variables. Equality (8.3) asserts that the “norm” of φ involves only the variables $(x, \xi) \in \mathbf{T}^2$. The definition (8.1) indicates that $\varphi(x, \xi)$ should be 1-periodic in ξ . With respect to the variable x we have the easily established property

$$\varphi(x+j, \xi) = e^{2\pi i j \xi} \varphi(x, \xi) \quad (8.4)$$

for each $j \in \mathbf{Z}$. The property (8.4) tells us how $\varphi(x, \xi)$, for $x \in [-\frac{1}{2}, \frac{1}{2}]$ and $\xi \in \mathbf{R}$, extends to all $x \in \mathbf{R}$. This shows that $|\varphi(x, \xi)|$ is 1-periodic in each of the variables $x, \xi \in \mathbf{R}$.

It is also easy to show that Z maps $L^2(\mathbf{R})$ onto $L^2(\mathbf{T}^2)$. Let $\varphi \in L^2(\mathbf{T}^2)$. Then $\int_{-\frac{1}{2}}^{\frac{1}{2}} |\varphi(x, \xi)|^2 d\xi < \infty$ for a.e. $x \in [-\frac{1}{2}, \frac{1}{2}]$. For each such x , φ , as a function of ξ , is a member of $L^2(\mathbf{T})$; thus, for a.e. $x \in [-\frac{1}{2}, \frac{1}{2}]$, $\varphi(x, \xi)$ has a Fourier series

$$\varphi(x, \xi) \sim \sum_{k \in \mathbf{Z}} c_k(x) e^{-2\pi i k \xi}$$

such that $\sum_{k \in \mathbf{Z}} |c_k(x)|^2 = \int_{-\frac{1}{2}}^{\frac{1}{2}} |\varphi(x, \xi)|^2 d\xi < \infty$. We then define a function f on \mathbf{R} by letting $f(x+k) = c_k(x)$ for each $k \in \mathbf{Z}$ and these $x \in [-\frac{1}{2}, \frac{1}{2})$. This a.e. defined function on \mathbf{R} belongs to $L^2(\mathbf{R})$ since

$$\int_{-\infty}^{\infty} |f(x)|^2 dx = \sum_{k \in \mathbf{Z}} \int_{-\frac{1}{2}}^{\frac{1}{2}} |c_k(x)|^2 dx = \int_{-\frac{1}{2}}^{\frac{1}{2}} \int_{-\frac{1}{2}}^{\frac{1}{2}} |\varphi(x, \xi)|^2 d\xi dx < \infty.$$

This shows that $\varphi(x, \xi) = (Zf)(x, \xi)$ for $(x, \xi) \in \mathbf{T}^2$. We then extend $\varphi(x, \xi)$ for all $x, \xi \in \mathbf{R}$ by making it 1-periodic in ξ and using (8.4) when $x \in [-\frac{1}{2}, \frac{1}{2})$ and $j \in \mathbf{Z}$.

We have, in fact, obtained the space M of all those measurable functions $\varphi(x, \xi)$, $(x, \xi) \in \mathbf{R}^2$, that are the general $\varphi \in L^2(\mathbf{T}^2)$ when (x, ξ) is restricted to \mathbf{T}^2 , are 1-periodic in ξ and satisfy (8.4). Moreover, if we introduce the norm

$$\|\varphi\|_M \equiv \left(\int_{\mathbf{T}^2} |\varphi(x, \xi)|^2 d\xi dx \right)^{\frac{1}{2}}$$

on M we have obtained

Theorem 8.1. *The linear operator Z maps $L^2(\mathbf{R})$ isometrically onto M .*

We now introduce the space \tilde{M} , a “companion” space to M . Essentially, \tilde{M} is obtained from M by reversing the roles of the variables x and ξ . More precisely, \tilde{M} consists of all functions $\tilde{\varphi}(x, \xi)$ on \mathbf{R}^2 that belong to $L^2(\mathbf{T}^2)$ when (x, ξ) is restricted to \mathbf{T}^2 , $\tilde{\varphi}$ is 1-periodic in x and its value for the general $\xi \in \mathbf{R}$ is given by equality

$$\tilde{\varphi}(x, \xi + \ell) = e^{-2\pi i \ell x} \tilde{\varphi}(x, \xi). \quad (8.5)$$

The norm of $\tilde{\varphi} \in \tilde{M}$ is, again, obtained from $L^2(\mathbf{T}^2)$:

$$\|\tilde{\varphi}\|_{\tilde{M}} = \left(\int_{\mathbf{T}^2} |\tilde{\varphi}(x, \xi)|^2 d\xi dx \right)^{\frac{1}{2}}.$$

There is a simple unitary map U from M onto \tilde{M} (these two spaces are, clearly, Hilbert spaces):

$$(U\varphi)(x, \xi) = e^{-2\pi i x \xi} \varphi(x, \xi) \equiv \tilde{\varphi}(x, \xi). \quad (8.6)$$

Let us explain why U maps M onto \tilde{M} . The general element, φ , of M , we have seen, is the Zak transform of an $f \in L^2(\mathbf{R})$: $\varphi(x, \xi) = (Zf)(x, \xi) = \sum_{k \in \mathbf{Z}} f(x+k) e^{-2\pi i k \xi}$. Thus,

$$(U\varphi)(x, \xi) = e^{-2\pi i x \xi} \sum_{k \in \mathbf{Z}} f(x+k) e^{-2\pi i k \xi} = \sum_{k \in \mathbf{Z}} f(x+k) e^{-2\pi i (x+k) \xi}. \quad (8.7)$$

We see, therefore, that $\tilde{\varphi} = U\varphi$ is 1-periodic in x . Since $\varphi(x, \xi), (x, \xi) \in \mathbf{T}^2$, is the general function of $L^2(\mathbf{T}^2)$, it is clear that $\tilde{\varphi}(x, \xi) = e^{-2\pi i x \xi} \varphi(x, \xi)$ is,

also, the general function of $L^2(\mathbf{T}^2)$. To see that $\tilde{\varphi}$ satisfies (8.5) we can use some of the ideas that showed Z was onto. Since $\tilde{\varphi} \in L^2(\mathbf{T}^2)$ we see that $\tilde{\varphi}(x, \xi) \cong \sum_{\ell \in \mathbf{Z}} c_\ell(\xi) e^{2\pi i \ell x}$ for a.e. $\xi \in [-\frac{1}{2}, \frac{1}{2}]$ and, for such a ξ , $\sum_{\ell \in \mathbf{Z}} |c_\ell(\xi)|^2 = \int_{-\frac{1}{2}}^{\frac{1}{2}} |\tilde{\varphi}(x, \xi)|^2 dx < \infty$. Define $g \in L^2(\mathbf{R})$ by letting $g(\xi + \ell) = c_\ell(\xi)$ for $\xi \in [-\frac{1}{2}, \frac{1}{2}]$ obtaining

$$\tilde{\varphi}(x, \xi) \cong \sum_{\ell \in \mathbf{Z}} g(\xi + \ell) e^{2\pi i \ell x}, \quad (8.8)$$

where “ \cong ” denotes the fact that the series on the right of (8.8) is the Fourier series of $\tilde{\varphi}(x, \xi)$ as a function of x for each of the $\xi \in [-\frac{1}{2}, \frac{1}{2}]$ described above. This enables us to define $\tilde{\varphi}(x, \xi)$ for a.e. $\xi \in \mathbf{R}$ by using (8.8):

$$\begin{aligned} \tilde{\varphi}(x, \xi + j) &= \sum_{\ell \in \mathbf{Z}} g(\xi + j + \ell) e^{2\pi i \ell x} \\ &= e^{-2\pi i j x} \sum_{\ell \in \mathbf{Z}} g(\xi + j + \ell) e^{2\pi i x(j + \ell)} = e^{-2\pi i j x} \tilde{\varphi}(x, \xi). \end{aligned}$$

This shows (8.5) and our definition of the space \tilde{M} is complete once we define $\|\tilde{\varphi}\|_{\tilde{M}}$ to be

$$\left(\int_{\mathbf{T}^2} |\tilde{\varphi}(x, \xi)|^2 dx d\xi \right)^{\frac{1}{2}}.$$

Observe that we have also shown that \tilde{M} is the image of an isometric Zak-like transform \tilde{Z} :

$$(\tilde{Z}g)(x, \xi) = \sum_{\ell \in \mathbf{Z}} g(\xi + \ell) e^{2\pi i \ell x} \equiv \tilde{\varphi}(x, \xi) \quad (8.9)$$

for $g \in L^2(\mathbf{R})$.

It is also natural to ask what is the relation between the functions f and g that satisfy

$$UZf = \tilde{Z}g. \quad (8.10)$$

We can determine this relation easily once we determine the inverse operators Z^{-1} and \tilde{Z}^{-1} . In fact, we have

$$\begin{aligned} (a) \text{ if } \varphi(x, \xi) \in M, \text{ then } (Z^{-1}\varphi)(x) &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \varphi(x, \xi) d\xi; \\ (b) \text{ if } \tilde{\varphi}(x, \xi) \in \tilde{M}, \text{ then } (\tilde{Z}^{-1}\tilde{\varphi})(\xi) &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \tilde{\varphi}(x, \xi) dx. \end{aligned} \quad (8.11)$$

To see this, recall that if $\varphi \in M$, Theorem 8.1 tells us that $\varphi(x, \xi) = \sum_{k \in \mathbf{Z}} f(x + \xi) e^{-2\pi i k \xi}$ for a unique $f \in L^2(\mathbf{R})$ and, as we have seen in (8.2), the sequence $\{f(x + k)\}$, $k \in \mathbf{Z}$, is, for a.e. $x \in \mathbf{Z}$, the sequence of Fourier coefficients of the $L^2(\mathbf{T})$

function $\varphi(x, \xi)$ that is 1-periodic in ξ for each such x . Hence, $\int_{-\frac{1}{2}}^{\frac{1}{2}} \varphi(x, \xi) d\xi = f(x)$, the zero coefficient of the function $\xi \rightarrow \varphi(x, \xi)$. This establishes (a) and the same argument, adapted to \tilde{M} , gives us (b).

The relation between f and g is then easily seen to be: $g = \hat{f}$ and $f = \check{g}$:

Theorem 8.2. (a) $\tilde{Z}^{-1}UZf = \hat{f} \equiv \mathcal{F}f$, the Fourier transform of f ;
(b) $Z^{-1}U^*\tilde{Z}g = \check{g} \equiv \mathcal{F}^{-1}g$, the inverse Fourier transform of g .

Proof. Assume $f \in L^1(\mathbf{R}) \cap L^2(\mathbf{R})$. From (8.7) we have $\tilde{\varphi}(x, \xi) = (U\varphi)(x, \xi) = \sum_{k \in \mathbf{Z}} f(x+k)e^{-2\pi i(x+k)\xi}$. Thus, applying \tilde{Z}^{-1} to the first and third expression in this equality we obtain (by (8.11a)):

$$(\tilde{Z}^{-1}UZf)(\xi) = \int_{-\frac{1}{2}}^{\frac{1}{2}} \sum_{k \in \mathbf{Z}} f(x+k)e^{-2\pi i(x+k)\xi} dx = \int_{-\infty}^{\infty} f(x)e^{-2\pi i x \xi} dx = \hat{f}(\xi).$$

This shows that $\tilde{Z}^{-1}UZ$ maps the space $L^1(\mathbf{R}) \cap L^2(\mathbf{R})$ onto the space of Fourier transforms of these functions. Since $L^1(\mathbf{R}) \cap L^2(\mathbf{R})$ is dense in $L^2(\mathbf{R})$, $\tilde{Z}^{-1}UZ$ is the unique extension of the map $f \rightarrow \hat{f}$ to a unitary operator, \mathcal{F} , on $L^2(\mathbf{R})$. Of course, \mathcal{F} is the Fourier transform on $L^2(\mathbf{R})$.

We also have that the inverse Fourier transform, \mathcal{F}^{-1} , satisfies $\mathcal{F}^{-1} = Z^{-1}U^*\tilde{Z}$. \square

This rather simple derivation of these important properties of the Fourier transform (in particular, that \mathcal{F} is unitary) has surprised some of our colleagues. That the Zak transform can be used to obtain these properties, however, is not new. Let us discuss the relevant history of this matter. Zak introduced the transform we denoted by \tilde{Z} in 1967 [7]. This operator, however, was also introduced by Gelfand [1] 17 years earlier; he included in this article an argument very similar to the one we gave above that showed the unitary property of the Fourier transform. Gelfand, actually, gives credit to A. Weil for this proof in a book [6] that came out essentially at the same time (this book is a “reprint” of an earlier one). In fact some authors call the operator \tilde{Z} the “Weil–Brezin” mapping; however, this occurs quite a bit later.

The two transforms we introduced, Z and \tilde{Z} , we believe, are rather natural and explain, perhaps, more clearly, the relationship between the Fourier transform (and its inverse) and the operators we introduced. The relations that were derived by the authors we quoted appear to us to involve a “Deus ex machina” in which the integral $\int_{-\frac{1}{2}}^{\frac{1}{2}} dx$ is applied to the expression $\sum_{k \in \mathbf{Z}} f(x+k)e^{-2\pi i(x+k)\xi}$ (see (8.7)) to obtain the Fourier transform. Our Theorem 8.2 does explain why this is done; the works we cite do not motivate the use of this integral for this purpose. If $h \in L^2(\mathbf{T})$ and $\sum_{n \in \mathbf{Z}} \hat{h}(n)e^{2\pi i n \xi}$ is the Fourier series of h we show how the well known equality $\int_{-\frac{1}{2}}^{\frac{1}{2}} |h(\xi)|^2 d\xi = \sum_{n \in \mathbf{Z}} |\hat{h}(n)|^2$ easily implies the Plancherel property of the Fourier transform \mathcal{F} . Our motivation for using the expression $e^{-2\pi i k \xi}$ (instead of $e^{2\pi i k \xi}$) in the definition of Z (see (8.1)) is that it leads us to the Fourier transform identities, see Theorem 8.2, more directly.

The spaces M and \tilde{M} are really quite interesting and, as we will show in another paper, are very much worth studying.

The operators Z, \tilde{Z}, U and their inverses play an important role in various different settings. In the next section, we shall present some observations that explain why this is the case.

8.2 More Uses of the Zak Transforms and Their Extensions

We first make some simple observations that follow immediately from the equalities in Theorem 8.2. Suppose f is, say, a Schwartz function (though we do not need so much smoothness and rapid vanishing at ∞); this allows us to formulate the result in Theorem 8.2 as a pointwise result valid for all $(x, \xi) \in \mathbf{R}^2$:

$$\sum_{k \in \mathbf{Z}} f(x+k) e^{-2\pi i(x+k)\xi} = \sum_{k \in \mathbf{Z}} \hat{f}(\xi+k) e^{2\pi i k x}. \quad (8.12)$$

If we set $x = 0 = \xi$, we obtain the *Poisson summation formula*

$$\sum_{k \in \mathbf{Z}} f(k) = \sum_{k \in \mathbf{Z}} \hat{f}(k). \quad (8.13)$$

Suppose we only set $x = 0$. We then have the equality

$$\sum_{k \in \mathbf{Z}} f(k) e^{-2\pi i k \xi} = \sum_{k \in \mathbf{Z}} \hat{f}(\xi+k). \quad (8.14)$$

If we make the additional band-limited assumption that $\text{Supp} \hat{f} \subset [-\frac{1}{2}, \frac{1}{2}]$ we then obtain

$$\hat{f}(\xi) = \sum_{k \in \mathbf{Z}} f(k) e^{-2\pi i k \xi}$$

for $\xi \in [-\frac{1}{2}, \frac{1}{2}]$. Let us multiply each side of this equality by $e^{2\pi i x \xi}$ and integrate both sides with respect to ξ . We obtain

$$f(x) = \int_{-\infty}^{\infty} \hat{f}(\xi) e^{2\pi i x \xi} d\xi = \sum_{k \in \mathbf{Z}} f(k) \int_{-\infty}^{\infty} \chi_{[-\frac{1}{2}, \frac{1}{2}]}(\xi) e^{2\pi i \xi(x-k)} d\xi$$

(we inserted $\chi_{[-\frac{1}{2}, \frac{1}{2}]}$ since $\text{Supp} \hat{f} \subset [-\frac{1}{2}, \frac{1}{2}]$). Since $\check{\chi}_{[-\frac{1}{2}, \frac{1}{2}]}(x) = \text{sinc}(x) \equiv \frac{\sin \pi x}{\pi x}$, we have

$$f(x) = \sum_{k \in \mathbf{Z}} f(k) \text{sinc}(x-k). \quad (8.15)$$

It is easy to check that the convergence of this series is absolute and in the norm of all the $L^p(\mathbf{R})$ spaces, $2 \leq p \leq \infty$. The reader will recognize this result as the *shape Shannon Sampling Theorem*. It tells us that if f is band limited, with $\hat{f}(\xi) = 0$ if $|\xi| > \frac{1}{2}$, then it is completely determined by its values on the integers (that is, by “sampling” f on the set \mathbf{Z}). There is a considerable literature devoted to extending this sampling result. The Zak transforms play an important role for obtaining these extensions; we intend to explain these results in future publications. We presented these facts as an illustration of the many uses of these operators in various areas of analysis.

It is clear that the two Zak transforms can be introduced in the setting of $L^2(\mathbf{R}^n)$, $n \geq 1$, and the results we derived in Sect. 8.1 can be extended in this more general case. These operators can also be considered to act on other $L^p(\mathbf{R}^n)$ spaces, $1 \leq p \leq 2$. In fact, the Zak transforms can be extended to much more general settings in which the groups \mathbf{Z}^n and \mathbf{T} are replaced by a locally compact abelian (LCA) group G and its dual \hat{G} . The space $L^2(\mathbf{R}^n)$ corresponds to a separable Hilbert space \mathbf{H} on which acts a unitary representation of G (see [3] and [4]). In these two articles just cited we study “Principal Shift Invariant Subspaces” of $L^2(\mathbf{R}^n)$, “Gabor Spaces” and their extensions in the LCA group setting. The extensions of the Zak transforms to these situations are also most useful in this context. Some of these extensions are also obtained by Gröchenig [2].

Acknowledgements The research of E. Hernández is supported by grants MTM2007 – 60952 of Spain and SIMUMAT S-0505/ESP-0158 of the Madrid Community Region. The research of H. Šikić, G. Weiss and E. Wilson is supported by the US-Croatian grant NSF-INT-0245238. The research of H. Šikić is also supported by the MZOS grant 037 – 0372790 – 2799 of the Republic of Croatia.

References

1. I. Gelfand, *Eigenfunction expansions for an equation with periodic coefficients*, Dokl. Akad. Nauk. SSSR, **76** (1950), pp. 1117-1120 (in Russian).
2. K. Gröchenig, *Aspects of Gabor analysis on locally compact abelian groups*, in: Gabor Analysis and Algorithms, H. G. Feichtinger and E. Strohmer, Editors. Birkhäuser, (2001), pp. 211-229.
3. E. Hernández, H. Šikić, G. Weiss, E. N. Wilson, *On the properties of the integer translates of a square integrable function*, Contemp. Math. **505**, (2010), pp. 233-249.
4. E. Hernández, H. Šikić, G. Weiss, E. N. Wilson, *Cyclic subspaces for unitary representations of LCA groups; generalized Zak transforms*, to appear in a volume dedicated to Andrzej Hulanicki publ. by Coll. Math in February (2010).
5. A. J. E. M. Janssen, *The Zak transform: a signal transform for sampled time-continuous signals*, Philips J. Research, **43** (1988), pp. 23-69.
6. A. Weil, *L'intégration dans les groupes topologiques et ses applications*, Hermann (Paris) (1951).
7. J. Zak, *Finite translations in solid state physics*, Phys. Rev. Lett. **19** (1967), pp. 1385-1397.

Part II

The Geometry of Large Data Sets

“Over the last twenty years we have seen the introduction of adaptive computational analytic tools that enable flexible transcriptions of the physical world. These tools enable orchestration of signals into constituents (mathematical musical scores) and opened doors to a variety of digital implementations/applications in engineering and science.

Of course I am referring to wavelet and various versions of computational Harmonic Analysis. The main concepts underlying these ideas involved the adaptation of Fourier analysis to changing geometries as well as multiscale structures of natural data. As such, these methodologies seem to be limited to analyze and process physical data alone.

Remarkably, the last few years have seen an explosion of activity in machine learning, data analysis and search, implying that similar ideas and concepts, inspired by signal processing might carry as much power in the context of the orchestration of massive high dimensional data sets. This digital data; text documents, medical records, music, sensor data, financial data etc. can be structured into geometries that result in new organizations of language and knowledge building. In these structures, the conventional hierarchical ontology building paradigm, merges with a blend of Harmonic Analysis and combinatorial geometry.

Conceptually these tools enable the integration of local association models into global structures in much the same way that calculus enables the recovery of a global function from a local linear model for its variation. As it turns out, such extensions of differential calculus into the digital data fields are now possible and open the door to the usage of mathematics similar in scope to the Newtonian revolution in the physical sciences. Specifically we see these tools as engendering the field of mathematical learning in which raw data viewed as clouds of points in high dimensional parameter space is organized geometrically much the same way as in our memory, simultaneously organizing and linking associated events, as well as building a descriptive language.

It is quite obvious to me that we can now attempt to translate some of the concepts and language of psychologists or neurologists into effective computational mathematical constructs, (not necessarily biological) performing similar functions. As an example, eigenvectors of an association matrix can provide a global Gestalt map of relations. While eigenvectors are mathematical constructs, a “wavelet paradigm”

translation can convert this approach to conventional hierarchical ontologies and language building. The main point for this kind of mathematical approach to the geometry and analysis of data is that it is an ab initio totally agnostic methodology which does not assume an a priori model or language. As such it can be used to validate our understanding, and more importantly to provide knowledge where the complexity exceeds our ability to model."

Raphy Coifman

Chapter 9

Harmonic Analysis of Digital Data Bases

Ronald R. Coifman and Matan Gavish

Abstract Digital databases can be represented by matrices, where rows (say) correspond to numerical sensors readings, or features, and columns correspond to data points. Recent data analysis methods describe the local geometry of the data points using a weighted affinity graph, whose vertices correspond to data points. We consider two geometries, or graphs – one on the rows and one on the columns, such that the data matrix is smooth with respect to the “tensor product” of the two geometries. This is achieved by an iterative procedure that constructs a multiscale partition tree on each graph. We use the recently introduced notion of Haar-like bases induced by the trees to obtain Tensor-Haar-like bases for the space of matrices, and show that an ℓ_p entropy conditions on the expansion coefficients of the database, viewed as a function on the product of the geometries, imply both smoothness and efficient reconstruction. We apply this methodology to analyze, de-noise and compress a term-document database. We use the same methodology to compress matrices of potential operators of unknown charge distribution geometries and to organize Laplacian eigenvectors, where the data matrix is the “expansion in Laplace eigenvectors” operator.

9.1 Introduction

“We should seek out unfamiliar summaries of observational material, and establish their useful properties.”
– J. W. Tukey [21]

There is a fruitful interplay between Harmonic analysis and the vast data analysis challenges that the scientific community is facing. Ideas that in the past have been applied for the analysis and extraction of information from physical systems are

M. Gavish (✉)

Department of Statistics, Stanford University, Sequoia Hall, Stanford, CA 94305, USA

e-mail: gavish@stanford.edu

and

Program in Applied Mathematics, Yale University, New Haven, CT 06511, USA

being increasingly applied, in their computational reincarnation, to organize and to extract information from high dimensional digital data sets of arbitrary source. Some examples are Laplacian eigenfunctions [2, 6] and wavelet bases [7].

In [20], J.O. Strömberg showed that the tensor product of Haar bases is extremely efficient in representing functions on product structures. For example, when $f : [0, 1]^d \rightarrow \mathbb{R}$ has bounded mixed derivatives, he considered the tensor product of Haar bases in all d dimensions and showed that f can be approximated with L_∞ error $\mathcal{O}(\varepsilon \log^{d-1}(\frac{1}{\varepsilon}))$ by shrinking to zero coefficients of the basis functions with support smaller than ε . As only $\mathcal{O}(\frac{1}{\varepsilon} \log^{d-1}(\frac{1}{\varepsilon}))$ coefficients survive, this yields a remarkable compression scheme.

In fact, *product structures* are among the most common data structures in existence. Consider the usual data matrix, abundant in multivariate statistics. By this we mean a rectangular array with n columns, representing (say) observations or individuals or data points, and p rows, representing attributes or variables measured on each data point. More generally, when each data point responds simultaneously to $d - 1$ variable sets, the data is given as a rank- d tensor. While the theory we will develop will include the general case of tensor data structures of arbitrary rank d , the examples discussed are data matrices where $d = 2$. Data matrices arise in text term-document analysis, microarray data analysis, gene association studies, sensor networks, recommendation engines, hyperspectral image processing – to name just a few scenarios. Here, the basic assumption underlying matrix analysis in multivariate statistics, namely that observations are independent and identically distributed, breaks down, since in general correlations exist among both rows and columns of the data matrix. Models for data of this form have been introduced in the statistical literature (e.g. Plaid Models [15] and Transposable Regularized Covariance Models [1]), yet the contrast between the overwhelming wealth of applications and the tools available remains considerable. In this work, we suggest a harmonic analysis approach that leads to a nonparametric model for data tensors, and in particular for data matrices.

Since correlations exist among both rows and columns of the data matrix, there is no longer a preferred dimension among $\{\text{rows, columns}\}$. Treating rows and columns on equal footing naturally leads to tensor analysis. Strömberg’s ideas, namely that tensor product of Haar bases sparsify smooth functions on a product space and that coefficients should be sorted by basis function support size, appears promising for efficiently analyzing matrices and higher order data structures. However, if we are to use this approach to analyze a matrix (a function on an abstract product space $M : \{\text{observations}\} \times \{\text{variables}\} \rightarrow \mathbb{R}$), we must first (1) make sense of the notion of a *geometry* and *Haar basis* on an abstract set such as $\{\text{observations}\}$ rather than on $[0, 1]$, (2) find an algorithm to construct two such geometries, relative to which a given data matrix is “regular”, and (3) find a computationally reasonable way to measure the regularity of a matrix with respect to two given geometries – on $\{\text{observations}\}$ and on $\{\text{variables}\}$. Only then can we hope to obtain an approximation result similar to Strömberg’s.

9.1.1 Outline

Our goal here is to follow this program. For (1), the local geometry of each set is described by a weighted affinity graph. For (2), in Sect. 9.2.1 we describe a straightforward procedure to integrate an affinity graph into a multiscale partition tree. In Sect. 9.2.2 we then describe an iterative procedure to construct two coupled geometries, namely affinity graphs, given a data matrix.

As an interesting example of organization, in Sect. 9.4 we consider the matrix of Laplacian eigenfunctions on a graph, where the entries represent the value of an eigenfunction at a point. As any two eigenvectors are orthogonal, their mutual distance is constant and a naive attempt to build an eigenvector organization fails. On the other hand, if we organize the graph into a contextual folder hierarchy and relate two eigenvectors also through their correlation on the folders at different scales, we obtain a dual organization of eigenvectors. In the case of a d -dimensional torus \mathbb{T}^d , using this procedure one recovers the usual Euclidean geometry of the dual group, namely the lattice \mathbb{Z}^d .

In Sect. 9.3, we consider a different operator, namely a potential operators on clouds of unknown geometries. We show that it can be organized to reveal the intrinsic geometry of their domain and range.

In Sect. 9.5, we turn to function bases. We build on [12, 13] where, in collaboration with B. Nadler, we showed that a partition tree on an abstract set X induces a “Haar-like” orthonormal basis for $\{f : X \rightarrow \mathbb{R}\}$ with wavelet basis properties. Thus, the construction of a coupled geometry on $\{\text{observations}\}$ and $\{\text{variables}\}$ also induces Haar-like bases, say $\{\psi_i\}$ and $\{\varphi_j\}$ respectively. We prove an approximation result of Strömberg’s type and show that a condition on the ℓ_p sum of expansion coefficients $\sum_{i,j} |\langle M, \psi_i \otimes \varphi_j \rangle|^p$ implies both efficiency of reconstruction and “bi-smoothness” of the data matrix with respect to the two trees constructed. Indeed the ℓ_p condition is better suited to this general scenario than a condition on mixed derivatives, which depends on the dimension d . In particular, this means that the ℓ_p condition measures the compatibility of the data matrix with the coupled geometry constructed, thus accomplishing (3) and yielding a stopping condition for the iterative procedure of Sect. 9.2.2. This approximation result leads to a decomposition of the dataset to a “smooth” part and an “irregular” part, the latter having a small support, in the spirit of the classical Calderón–Zygmund decomposition. In Sect. 9.7, we describe how a tensor-Haar basis, induced by a coupled geometry that is compatible with the data matrix, can be used to perform compression and statistical tasks such as denoising and missing value imputation on the data matrix.

9.1.2 An Example

In Sect. 9.7, we study a term-document data set: a matrix M whose entry $M_{i,j}$ is the occurrence frequency of word i in document j . We define a graph structure on the documents, in which we link documents with a preponderance of common




Fig. 9.1 The Science News data matrix discussed in Sect. 9.7.2. The partition trees on rows and columns are shown. The rows and columns have been re-ordered by the trees depth-first order

vocabulary. This particular graph is crude: highly related documents, which should be nearest neighbors in this graph, share just a small set of highly correlated words – but this correlation may drown in the overall noise. Still this graph allows to organize documents into an inaccurate hierarchy of contextual folders. In order to obtain an improved hierarchy, it is useful to consider conceptual groupings of words, thereby leading us to build an analogous graph structure on the vocabulary, in which we link two words if they occur frequently in the same documents as well as in the contextual folders of documents. This leads to a hierarchy of conceptual word folders, which are then used to refine and rebuild the original document graph and contextual folder structure, and so on (Fig. 9.1).

In each stage of the iteration, the partition trees constructed on the data set induce two Haar-like bases, one on documents and one on words. The original data base can now be viewed as a function on the tensor product of the two graphs, and can therefore be expanded in the tensor product of the two bases. At each stage we compute the ℓ_1 entropy of the tensor-Haar coefficients as way of measuring the fit between the geometry and data base, and stop the iterations when no improvement is detected. Observe that by construction this procedure is invariant under permutations of matrix rows and columns, providing a tool for recovering an intrinsic organization of both – even when we are given a matrix with no known order on its rows or columns. Once a coupled geometry that is compatible with the matrix is achieved, we expand the data matrix in the tensor-Haar basis and process the data

base in the coefficient domain. For instance, the data matrix can now be compressed or de-noised using standard signal processing methodologies. While the Haar bases are not smooth, their construction is random and can be repeated and averaged, in order to eliminate some of the artifacts created, following [5].

In Sect. 9.8, we return to the example of potential operator and show that by transforming to the tensor-Haar coefficient matrix, the operator becomes extremely compressed. In other words, by treating a potential operator as a data matrix, recovering its coupled geometry, transforming to a tensor-Haar basis and thresholding coefficients, we get a fast numerical scheme.

Note that the ℓ_1 entropy condition of coefficients of an orthogonal expansion is used to quantify the “sparseness” of the expansion – a well known notion in the Compressed Sensing literature [4, 11]. It is an easy observation that a function can be recovered to mean square error ε times the ℓ_1 entropy by using only the coefficients larger than ε , whose number does not exceed the ℓ_1 entropy times $\frac{1}{\varepsilon}$. In the special case of the Haar or tensor-Haar expansions, *only coefficients of Haar functions with support volume larger than ε are required*. This conveniently eliminates the need to sort the large set of coefficients.

9.2 The Coupled Geometry of Questionnaires and Databases

Suppose that

$$M = \begin{pmatrix} | & & | \\ v_1 & \cdots & v_n \\ | & & | \end{pmatrix}$$

is a data matrix, whose columns correspond to observations. In order to organize the observations as the vertices of an affinity graph, a kernel-based approach would be to take the affinity $W_{i,j}$ between observations i and j to be, for instance, $\exp\left(-\|v_i - v_j\|^2\right)$. In modern data analysis scenarios, the dimensionality of v_i is often comparable or larger than n . In this scenario, different sets of observations might respond to different sets of variables whereas the correlation between any two complete vectors v_i and v_j is either very large or very small. Consider for example automatic medical tests that measure multiple parameters such as chemical blood composition or brain activity. Different patient groups (corresponding to different medical situations) will have very different response to different variable sets. In these scenarios, quantities such as $\exp\left(-\|v_i - v_j\|^2\right)$ do not provide a useful affinity among observations. Instead, we should consider partitioning both the observation and variable sets.

9.2.1 Describing a Graph's Global Structure Using a Partition Tree

The main computational tool we will use to construct a *coupled geometry* and, later, the orthonormal bases they induce, is a *partition tree*. By this we mean a sequence of increasingly refined partitions, readily described by a tree. As an affinity graph only holds “local” geometrical information relating each point to its neighbors, it does not directly allow computations involving the large scale structure of the graph. A convenient way to achieve this is by “integrating” the local distances into a partition tree.

Let $X = \{x_1, \dots, x_N\}$ be a finite set. Consider a sequence of L finer and finer partitions of X , denoted $\mathcal{X}^1, \dots, \mathcal{X}^L$. For each $1 \leq \ell \leq L$, the partition at level ℓ is composed of $n(\ell)$ mutually disjoint sets, which we call *folders*,

$$\mathcal{X}^\ell = \{X_1^\ell, \dots, X_{n(\ell)}^\ell\} \quad (9.1)$$

such that

$$X = \bigsqcup_{k=1}^{n(\ell)} X_k^\ell. \quad (9.2)$$

The finest partition, at level $\ell = L$, is composed of $n(L) = N$ singleton folders: $X_k^L = \{x_k\}$ for $k = 1 \dots N$. The coarsest partition, at level $\ell = 1$, is composed of a single folder, $X_1^1 = X$. The partitions are *nested* in the sense that for $1 < \ell \leq L$, each folder $X_k^\ell \in \mathcal{X}^\ell$ is a subset of a folder from $\mathcal{X}^{\ell-1}$. We let $\text{subfolders}(\ell, k) \subset \{1 \dots n(\ell+1)\}$ be the indices such that

$$X_k^\ell = \bigsqcup_{j \in \text{subfolders}(\ell, k)} X_j^{\ell+1}. \quad (9.3)$$

There are many approaches in the literature for constructing a multiscale partition tree given a symmetric adjacency matrix W describing a weighted graph on the vertex set X . We suggest using the following random bottom-up construction. The most refined nontrivial partition, \mathcal{X}^{L-1} is created using a standard “k-means type” approach: For a fixed “radius” $\rho > 0$, a maximal set of centroid points $\{z_i\} \subset X$ such that $i \neq j \Rightarrow W(z_i, z_j) \geq \rho$ is selected at random. A partition of X is now obtained by grouping points by affinity from the centroids, namely setting $X_i^{L-1} = \{x \in X \mid W(x, z_i) > W(x, z_j) \forall i \neq j\}$ for $i = 1 \dots n(L-1)$, where $n(L-1) = \#\{z_i\}$.

In order to group the folder set $\mathcal{X}^{\ell-1} = \{X_1^{L-1} \dots X_{n(L-1)}^{L-1}\}$ and obtain the next partition, define an affinity between folders by

$$\tilde{W}(i, j) = \left\langle W \mathbf{1}_{X_i^{L-1}}, W \mathbf{1}_{X_j^{L-1}} \right\rangle = \sum_{x \in X_i^{L-1}} \sum_{y \in X_j^{L-1}} W^2(x, y),$$

so that the affinity between folders is measured in the next time-scale. The partition procedure is repeated for the set $\mathcal{X}^{\ell-1} = \{X_1^{L-1} \dots X_{n(L-1)}^{L-1}\}$ with the affinity \tilde{W} to yield the next partition \mathcal{X}^{L-2} , and so on until a trivial partition $X_1^1 = X$ is reached. Below, we refer to the partition sequence $\mathcal{X}^1, \dots, \mathcal{X}^L$ as a partition tree \mathcal{T} on X .

We remark that for data sets of size small enough to allow computation of graph Laplacian eigenvectors, it is sometimes preferable to embed the set X in Euclidean space first (e.g. using a Diffusion Embedding [6]) and construct the partition tree there.

9.2.2 An Iterative Procedure to Construct a Coupled Geometry

Suppose that $f, g : X \rightarrow \mathbb{R}$ are two functions on X and $\rho(\cdot, \cdot)$ is an affinity between functions. Typical choices for $\rho(f, g)$ include $\exp\left(-\frac{1}{\varepsilon} \sum_{x \in X} (f(x) - g(x))^2\right)$ for some $\varepsilon > 0$. Another example is the absolute value of correlation coefficient, $\left| \frac{\text{cov}(f, g)}{\sigma(f) \cdot \sigma(g)} \right|$, whenever it exceeds a given cutoff. If f and g orthogonal, $\rho(f, g)$ would hold no information. But it is possible that f and g are strongly correlated if restricted to part of X . A classical example is sine functions of close frequencies, when compared on a small subset of $[0, 1]$.

Given a partition tree \mathcal{T} on X , define the affinity $\rho_{\mathcal{T}}(f, g)$ as follows. Let $\tilde{f}_{\mathcal{T}}(X_k^{\ell}) = \frac{1}{|X_k^{\ell}|} \sum_{x \in X_k^{\ell}} f(x)$ denote the average of f on the folder X_k^{ℓ} of \mathcal{T} . In particular $\tilde{f}_{\mathcal{T}}(X_k^L) = f(x_k)$. Thus $\tilde{f}_{\mathcal{T}} : \mathcal{T} \rightarrow \mathbb{R}$ is an extension of f , augmenting its original values by its average values on all folders in all levels of \mathcal{T} . By setting $\rho_{\mathcal{T}}(f, g) = \rho(\tilde{f}_{\mathcal{T}}, \tilde{g}_{\mathcal{T}})$, we take into account the mutual behavior of f and g on all levels of \mathcal{T} .

Now consider a matrix M . Denote the set of columns of M by X , and the set of its rows by Y . When M is a data set, we interpret X as observations and Y as variables, features or sensors. Thus $M : X \times Y \rightarrow \mathbb{R}$.

Definition 9.1. For each $y \in Y$, the row $M_y(x) : x \mapsto M(x, y)$ of M is a function on X . For a given a partition tree \mathcal{T}_X on X , define the *dual affinity* on Y by $W_{\mathcal{T}_X}(y_1, y_2) = \rho_{\mathcal{T}_X}(M_{y_1}, M_{y_2})$. Similarly, for each $x \in X$, the column $M_x(y) : y \mapsto M(x, y)$ of M is a function on Y . For a given partition tree \mathcal{T}_Y on Y , define the *dual affinity* on X by $W_{\mathcal{T}_Y}(x_1, x_2) = \rho_{\mathcal{T}_Y}(M_x, M_{x_2})$.

We thus arrive at the following procedure to refine a given affinity W on the set X of columns of X :

- Algorithm 9.1.**
1. Integrate the affinity W as in Sect. 9.2.1 to obtain a partition tree \mathcal{T}_X on X .
 2. Compute the dual affinity $W_{\mathcal{T}_X}$ on Y as in Definition 9.1.
 3. Integrate the affinity $W_{\mathcal{T}_X}$ to obtain a partition tree \mathcal{T}_Y on Y .
 4. Compute the *refined version* \tilde{W} of W by setting \tilde{W} to the dual affinity $W_{\mathcal{T}_Y}$ on X .

By iterating this cycle, we obtain stable and refined affinities both on X and on Y . Given a matrix M , the initial affinity W on X is either taken from an external source or taken to be the affinity between the columns of M . This procedure was introduced in [22] and applied for automatic translation using term-document matrices of the same text in different languages.

We now consider two examples from mathematics, rather than data analysis, which illustrate the notion of *coupled geometry* and the possibilities offered by an ability to recover it.

The experiments we describe below suggest that this iterative procedure converges to affinities (W_X, W_Y) , which capture what we informally call *the coupled geometry of the matrix M* . However, we still need a quantitative way to determine convergence, and a stopping rule for the iterations. Equally important, we must be able to evaluate the performance of this algorithm and to compare it with other approaches to the same task. This leads to the following key questions:

1. How to quantify and measure the compatibility of the database to a proposed coupled geometry?
2. How can a coupled geometry, which is compatible with the database, be used for analyzing it?

These questions are answered by the theory developed in Sect. 9.6.

9.3 Example: Numerical Compression of a Potential Operator

As a first example of finding the *coupled geometry* of matrix rows and columns, consider the potential interaction between two point clouds in three dimension as in Fig. 9.2c. Let $\{x_i\}_{i=1}^N \subset \mathbb{R}^3$ and $\{y_j\}_{j=1}^N \subset \mathbb{R}^3$ denote the blue points (one dimensional helix) and the red points (two-dimensional sheet) respectively. The Coulomb potential operator is defined by

$$M_{i,j} = \frac{1}{\|x_i - y_j\|}.$$

Instead of the potential matrix M (Fig. 9.2a), we are denied the spatial layout of the point clouds and given only the potential matrix with rows and columns in random order.

In this case, after one iteration of the above procedure we recognize that there is a one-dimensional and a two-dimensional structure involved. Figure 9.3 shows a diffusion embedding into three dimensional space (as in [6]) of graphs obtained for the matrix columns, corresponding to the two-dimensional charge plate, and of the graph obtained for the matrix rows, corresponding to the one-dimensional charge helix.

In contrast with the fast multipole method [14], which would treat the points in three dimensions, here a point is placed in the overall geometry according to its

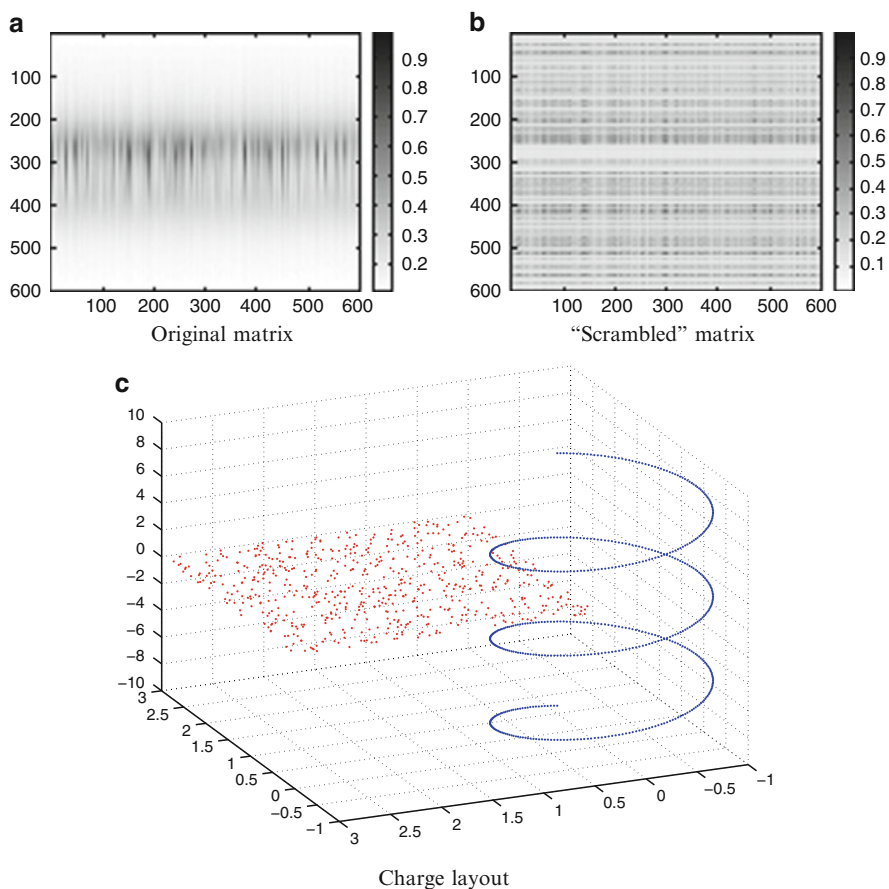


Fig. 9.2 (a) Original matrix of potential operator. (b) After permuting rows and columns. (c) The spatial layout: the helix consists of the points $\{x_i\}_{i=1}^{600}$ and the flat sheet consists of the points $\{y_i\}_{i=1}^{600}$

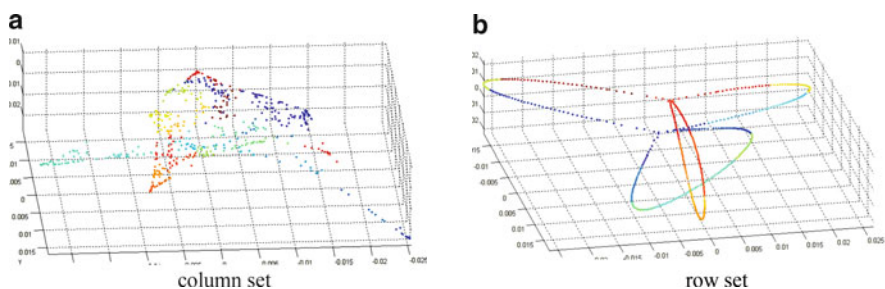


Fig. 9.3 The three-dimensional Diffusion embedding of the "scrambled" potential matrix (Sect. 9.3) row set and column set recovers their correct intrinsic dimensionality and geometry. Symmetries of the spatial layout are present in the potential operator and are reflected in this intrinsic geometry

interactions with other points alone. This suggests a general method for organizing (and, as we will see in Sect. 9.8, for compression and fast calculations) of quite general potential operators on unknown geometries.

9.4 Example: Recovering the Intrinsic Geometry of Laplacian Eigenfunctions

A second, interesting example for recovering a *coupled geometry* using the iterative procedure of Sect. 9.2 involves organizing a matrix whose rows are samples of some manifold Laplacian eigenfunctions.

On a general manifold, the only a priori relationship between Laplacian eigenfunctions is through the corresponding eigenvalues, which provide a degenerate one dimensional geometry on the eigenfunctions. However, it is often clear that the set of eigenfunctions carries a much richer structure. As a canonical example, consider the Laplacian eigenvectors on the d -dimensional torus \mathbb{T}^d . This set is indexed by the dual group, \mathbb{Z}^d . The structure of eigenfunctions can thus be organized using the Euclidean geometry \mathbb{Z}^d . This organization, whereby (in $d = 2$, say) the eigenfunctions $\phi_{k,m}(x,y) = \sin(2\pi kx) \sin(2\pi my)$ and $\phi_{k,m}(x,y) = \sin(2\pi(k+1)x) \sin(2\pi(m+1)y)$ seem *identical* on scales much larger and much smaller than $(\frac{1}{k}, \frac{1}{m})$ is fundamental in Fourier analysis. Extending this fact, we claim that by considering the correlation between eigenvectors on different levels of the a multiscale decomposition of the manifold, we can define a “dual geometry” on eigenfunctions. This introduces a variety of questions in analysis, concerning the relation between this dual geometry and the properties of the manifold.

9.4.1 Recovering the Geometry of Circle Eigenfunctions Using “Partition-Based Correlation”: A Calculation

Model the d -dimensional torus \mathbb{T}^d by $[0,1]^d$. The fixed dimension d is suppressed below. For $\ell \geq 0$, denote by $K_\ell = 2^{-\ell} \cdot \{0, \dots, 2^\ell - 1\}^d$ the square mesh of resolution $2^{-\ell}$ in \mathbb{T}^d . Consider the partition tree \mathcal{T} with partition at level ℓ given by $\mathcal{X}^\ell = \biguplus_{\mathbf{k} \in K_\ell} (\mathbf{k} + [0, 2^{-\ell})^d)$. Let ψ be a function such that $\sum_{\mathbf{k} \in K_\ell} \psi(2^\ell \mathbf{x} - \mathbf{k})$ approximates the constant function $\mathbf{1}$ on \mathbb{T}^d for each ℓ , $\int \psi(\mathbf{x}) d\mathbf{x} = 1$, $\psi \geq 0$, $\frac{\partial}{\partial \xi_i} \hat{\psi}(0) = 0$ and $\frac{\partial^2}{\partial \xi_i \partial \xi_j} \hat{\psi}(0) = C \delta_{i,j}$ where $\hat{\psi}$ is the Fourier transform. For example, ψ can be the indicator of $[0, 1)^d$.

Definition 9.2. Define the partition-based correlation (with respect to the partition \mathcal{X}^ℓ and weight ψ) of functions $f, g \in L_2(\mathbb{T}^d)$ by

$$\rho_\ell(f, g) = \sum_{\mathbf{k} \in K_\ell} \left| \int_{\mathbb{T}^d} f(\mathbf{x}) \overline{g(\mathbf{x})} \psi(2^\ell \mathbf{x} - \mathbf{k}) d\mathbf{x} \right|.$$

For eigenfunctions of \mathbb{T}^d corresponding to eigenvalues $\mathbf{m}, \mathbf{m}' \in \mathbb{Z}^d$ we have

$$\begin{aligned} \rho_\ell(e^{i\mathbf{m} \cdot \mathbf{x}}, e^{i\mathbf{m}' \cdot \mathbf{x}}) &= \sum_{\mathbf{k} \in K_\ell} \left| \int_{\mathbb{T}^d} e^{i(\mathbf{m} - \mathbf{m}') \cdot \mathbf{x}} \psi(2^\ell \mathbf{x} - \mathbf{k}) d\mathbf{x} \right| \\ &= \sum_{\mathbf{k} \in K_\ell} 2^{-\ell d} \left| \hat{\psi}\left(\frac{\mathbf{m} - \mathbf{m}'}{2^\ell}\right) \right| = \left| \hat{\psi}\left(\frac{\mathbf{m} - \mathbf{m}'}{2^\ell}\right) \right|. \end{aligned}$$

To see that this recovers the Euclidean affinity of \mathbb{Z}^d , recall that an affinity needs hold only for neighbors. Indeed, to second order in $2^{-\ell} \|\mathbf{m} - \mathbf{m}'\|$, this gives

$$\rho_\ell(e^{i\mathbf{m} \cdot \mathbf{x}}, e^{i\mathbf{m}' \cdot \mathbf{x}}) \approx \left| 1 - \frac{C}{2} \cdot 2^{-2\ell} \cdot \|\mathbf{m} - \mathbf{m}'\|^2 \right|$$

so that

$$2^{2\ell} \left(1 - \rho_\ell(e^{i\mathbf{m} \cdot \mathbf{x}}, e^{i\mathbf{m}' \cdot \mathbf{x}}) \right) \propto \|\mathbf{m} - \mathbf{m}'\|^2.$$

To improve the approximation, we can define $\rho(f, g) = \sum_{\ell \geq 1} 2^{-\ell} \rho_\ell(f, g)$, so that

$$\rho(e^{i\mathbf{m} \cdot \mathbf{x}}, e^{i\mathbf{m}' \cdot \mathbf{x}}) = \sum_{\ell \geq 1} 2^{-\ell} \rho_\ell(e^{i\mathbf{m} \cdot \mathbf{x}}, e^{i\mathbf{m}' \cdot \mathbf{x}}) \approx \left| 1 - \text{const} \cdot \|\mathbf{m} - \mathbf{m}'\|^2 \right|.$$

For example, taking $d = 1$ and $\psi = \mathbf{1}_{[0,1]}$ we get

$$\rho_\ell(e^{imx}, e^{im'x}) = \left| \text{sinc}\left((m - m')2^{-(\ell+1)}\right) \right| \approx \left| 1 - \frac{1}{6} \left((m - m')2^{-(\ell+1)} \right)^2 \right|.$$

9.4.2 Recovering the Geometry of Circle Eigenfunctions Using “Partition-Based Correlation”: A Computation

Given the above, it is tempting to ask whether one can use partition-based correlation, or similar function affinity notions that are based on a partition tree, to recover useful geometries on sets of manifolds and graphs eigenfunctions. However, it seems that there are very few cases where this question can be tackled analytically. The following simple experiment suggests that the iterative procedure of Sect. 9.2 enables one to study this question empirically.

In an experiment, summarized in Fig. 9.4 below, we recover the dual geometry of the eigenfunctions of the circle \mathbb{T}^1 . Let x_1, \dots, x_N be equally spaced points in $[0, 1]$ and consider the $2m$ by N matrix

$$M_{i,j} = \begin{cases} \sin(2\pi k) & k = 1 \dots m \\ \cos(2\pi k) & k = m + 1 \dots 2m \end{cases}$$

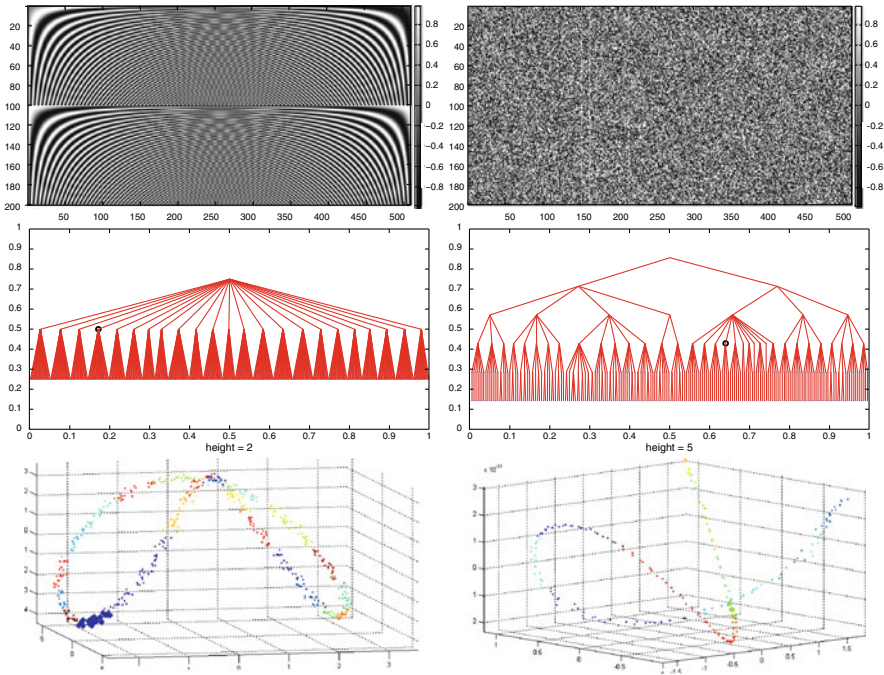


Fig. 9.4 Recovering the eigenfunction organization of \mathbb{T}^1 : an experiment for a sines–cosines matrix on $N = 512$ points and frequencies $1 \dots m = 100$ (see Sect. 9.4.2). *Top left*: the original matrix $M_{i,j}$. *Top right*: the matrix after a random permutation of the rows and a random permutation of the columns. *Middle left*: the partition tree on the points $\{x_i\}$ generated by the iterative procedure. Points in the same folder in the middle level have close x values. For example, points in the folder marked by a circle contain all the points in $[0.42, 0.49]$. *Middle right*: the partition tree on the functions generated by the iterative procedure. Functions in the same folder have similar frequencies. For example, the functions in the folder marked by a circle are $\sin(2\pi \cdot 81)$, $\sin(2\pi \cdot 82)$, $\cos(2\pi \cdot 80)$, $\cos(2\pi \cdot 81)$. *Bottom left*: A Diffusion Embedding visualization in 3-d space of the of the affinity graph on the points $\{x_i\}$, generated by the iterative procedure. The same points folder as in the middle left panel is highlighted. The curve recovers the geometry of the points $\{x_i\}$ along the circle \mathbb{R}/\mathbb{Z} exactly as points along the curve have increasing x value modulo 1. *Bottom right*: A Diffusion Embedding [6] visualization in 3-d space of the of the affinity graph on the functions, generated by the iterative procedure. The same points folder as in the middle right panel is highlighted. The curve recovers the geometry of the functions exactly as functions along the curve have increasing frequency.

The rows of M are orthogonal. However, eigenfunctions of similar frequencies behave similarly on folders of scale comparable to their frequency. The iterative procedure is designed to detect precisely this kind of similarity. By augmenting the values of each eigenfunctions by its values on folders of all scales and taking correlations along these augmented functions, we calculate a version of the correlation-based affinity and indeed recover the Euclidean affinity between the frequencies. To emphasize the organizational power of this procedure, we apply random permutations to the rows and columns of M before invoking the iterative procedure.

In a work in progress, we are applying the same procedure to investigate dual geometries of eigenfunctions in situations where an analytical approach is unknown, such the connected sum of two tori in \mathbb{R}^3 .

9.5 Tensor-Haar-Like Bases

We now proceed to answer the two questions that were stated at the end of Sect. 9.2. The main tool to evaluate the compatibility of a coupled geometry to a given data matrix, and later to process a data matrix, tensor-Haar-like bases. A Haar-like bases is a multiscale, localized orthonormal bases induced by a partition tree. These bases were introduced in [12, 13]. An application to semi-supervised learning is included in [12]. A similar construction with applications to pattern detection in networks appears in [18].

As with any hierarchical partition construction on an abstract space, this section is inevitably heavy with notation. In Sect. 9.6.5, we discuss the familiar Euclidean case using more or less the same notation.

9.5.1 Haar-Like Bases

Definition 9.3. Let \mathcal{T} be a partition tree (recall Sect. 9.2.1) with levels $\mathcal{X}^1, \dots, \mathcal{X}^L$ on a set X . We say that \mathcal{T} is $(\underline{B}, \overline{B})$ -balanced, or $(\underline{B}, \overline{B})$ -regular, if

$$\underline{B} \leq \frac{|X_j^{\ell+1}|}{|X_k^\ell|} \leq \overline{B} \quad (9.4)$$

for any $j \in \text{subfolders}(\ell, k)$ (recall (9.3)).

There is a discrete analog of multiresolution analysis associated with a partition tree. Let $V = \{f \mid f : X \rightarrow \mathbb{R}\}$. Each partition \mathcal{X}^ℓ induces a subspace $V^\ell \subset V$ by $V^\ell = \text{Span}_{\mathbb{R}} \left\{ \mathbf{1}_{X_1^\ell}, \dots, \mathbf{1}_{X_{n(\ell)}^\ell} \right\}$. As $V^\ell \subset V^{\ell+1}$, we write W^ℓ ($1 \leq \ell < L$) for the orthogonal complement of V^ℓ in $V^{\ell+1}$. Clearly $V^L = \left[\bigoplus_{\ell=1}^{L-1} W^\ell \right] \oplus V^1$.

Definition 9.4. 1. A Haar-like basis Ψ^ℓ for W^ℓ is an orthonormal basis of the form

$$\Psi^\ell = \bigcup_{k=1}^{n(\ell)} \left[\{\psi_{\ell,k,j}\}_{j=1}^{\#\text{subfolders}(\ell,k)-1} \right]$$

- where for each $1 \leq k \leq n(\ell)$, the function $\psi_{\ell,k,j}$ is supported on the folder X_k^ℓ .
2. A Haar-like basis Ψ for V is a union of Haar-like bases for each W^ℓ , $1 \leq \ell \leq L-1$, together with the normalized constant function on X , $\psi_0 \equiv \frac{1}{\sqrt{N}}$. Namely,

$$\Psi = \{\psi_0\} \cup \bigcup_{\ell=1}^{L-1} \mathfrak{B}^\ell = \{\psi_0\} \cup \bigcup_{\ell=1}^{L-1} \bigcup_{k=1}^{n(\ell)} \left[\{\psi_{\ell,k,j}\}_{j=1}^{\#\text{subfolders}(\ell,k)-1} \right].$$

Figure 9.5 illustrates a Haar-like basis induced by a partition tree on a small set.




Fig. 9.5 An illustration of a Haar-like basis on a set of 9 points

Remark. 1. Each basis function $\psi_{\ell,k,j}$ is associated with its support folder X_k^ℓ in the partition \mathcal{X}^ℓ . The number of basis functions associated to the same folder X_k^ℓ is $\#\text{subfolders}(\ell, k) - 1$.

2. These functions resemble the classical Haar functions in the following sense:

- Since $W^\ell \subset V^{\ell+1}$, each $\psi_{\ell,k,j}$ is piecewise constant on the folders of $\mathcal{X}^{\ell+1}$.
- Since $\psi_{\ell,k,j}$ is supported on the folder X_k^ℓ , it is nonzero only on these folders of $\mathcal{X}^{\ell+1}$ that are subfolders of X_k^ℓ . In other words, $\psi_{\ell,k,j}$ is a linear combination of $\left\{ \mathbf{1}_{X_i^{\ell+1}} \right\}_{i \in \text{subfolders}(\ell, k)}$.
- Since $W^\ell \perp V^\ell$, we have $\langle \psi_{\ell,k,j}, \mathbf{1}_{X_k^\ell} \rangle = 0$, so that $\psi_{\ell,k,j}$ is orthogonal to the constant function on X_k^ℓ .

9.5.2 Tensor Product of Haar-Like Bases

Now suppose that for each $\alpha = 1 \dots d$, $X[\alpha]$ is a set and that $\mathcal{T}[\alpha]$ is a partition tree on it with levels $\mathcal{X}[\alpha]^\ell = \{X[\alpha]_k^\ell\}_{k=1}^{n[\alpha](\ell)}$ for $\ell = 1 \dots L[\alpha]$. Consider the product $X = X[1] \times \dots \times X[d]$ with the normalized counting measure $|\cdot|$. If $\#X[\alpha]_k^\ell$ is the cardinality of $X[\alpha]_k^\ell$, then

$$\left| \prod_{\alpha=1}^d X[\alpha]_{k(\alpha)}^{\ell(\alpha)} \right| = \left(\prod_{\alpha=1}^d \#X[\alpha] \right)^{-1} \cdot \prod_{\alpha=1}^d \#X[\alpha]_{k(\alpha)}^{\ell(\alpha)}.$$

Definition 9.5. 1. Denote by

$$\mathcal{R} = \left\{ \prod_{\alpha=1}^d X[\alpha]_{k(\alpha)}^{\ell(\alpha)} \mid 1 \leq \ell(\alpha) \leq L[\alpha] \text{ and } 1 \leq k(\alpha) \leq n[\alpha](\ell) \right\}$$

the set of all “rectangles”, namely all products of folders of the partition trees $\mathcal{T}[1] \dots \mathcal{T}[d]$.

2. Write $x = (x_1, \dots, x_d)$ for an element of the product space $X = \prod_{\alpha=1}^d X[\alpha]$.
3. If $\Psi[\alpha] = \{\psi[\alpha]_{\ell,k,j}\}$ is a Haar-like basis induced by the partition tree $\mathcal{T}[\alpha]$ for $\alpha = 1 \dots d$, then

$$\Psi = \left\{ x \mapsto \prod_{\alpha=1}^d \psi[\alpha]_{\ell,k,j}(x_\alpha) \mid \psi[\alpha]_{\ell,k,j} \in \Psi[\alpha] \right\}$$

is the corresponding tensor-Haar-like basis. Clearly, it is an orthonormal basis for $\{f : X \rightarrow \mathbb{R}\}$.

4. Recall that each basis function $\psi[\alpha]_{\ell,k,j} \in \Psi[\alpha]$ is associated with its support folder $X[\alpha]_k^\ell$. We can thus write $R(\psi) \in \mathcal{R}$ for the support rectangle of a tensor basis function $\psi \in \Psi$.
5. It will be convenient to enumerate the elements of Ψ by $\Psi = \{x \mapsto \psi_q(x)\}_{q=1}^N$, where $N = \prod_{\alpha=1}^d \#X[\alpha]$.
6. Assume that for each α the tree $\mathcal{T}[\alpha]$ is $(\underline{B}[\alpha], \bar{B}[\alpha])$ – balanced (in the sense of Definition 9.3). Define $\tau(\Psi) = (\prod_{\alpha=1}^d \underline{B}[\alpha])^{-1}$.

To stress the fact that for a data matrix, computing a tensor-Haar-like basis function coefficient involves a non-trivial average of parts of the matrix, Fig. 9.6 shows an example from the Science News dataset discussed in Sect. 9.7. Top, a tensor-Haar-like function in color map drawn using the original organization of the data




Fig. 9.6 An example of a tensor-Haar-like function on the Science News data matrix. *Top:* In the original rows and columns order of the matrix. *Bottom:* After sorting rows and columns according to the trees which yielded this tensor-Haar-like function. In the original order of the rows and columns, namely the order in which the data matrix is observed, an inner product with this basis function is a highly nontrivial averaging operation

matrix. Bottom, the same function after reordering the rows and columns according to a in-order traversal of the partition trees, which is the ordering used to construct the basis functions.

We conclude this section with an estimate of $\int_X \psi$ for use in the approximation theorem below. Recall that the absolute value of a classical Haar function is constant. This gives $\|\psi\|_\infty^2 \cdot |R(\psi)| = 1$ and hence

$$\int_X |\psi|^p = \int_X |R(\psi)|^{-\frac{p}{2}} = |R(\psi)|^{1-\frac{p}{2}}$$

for a classical Haar function ψ . This result extends to our setting using the “balance” constant $\tau(\Psi)$ above.

Lemma 9.1. *Let Ψ be a tensor-Haar-like basis. Using the notation of Definition 9.5, for any $\psi \in \Psi$ and any $0 < p < 2$ we have*

$$\int_X |\psi|^p \leq \tau(\Psi)^{\frac{p}{2}} |R(\psi)|^{1-\frac{p}{2}}$$

Proof. We first show that

$$\max_{\psi \in \Psi} \{ \|\psi\|_\infty^2 \cdot |R(\psi)| \} \leq \tau(\Psi). \quad (9.5)$$

It is enough to show that for each $1 \leq \alpha \leq d$,

$$\max_{\psi[\alpha] \in \Psi[\alpha]} \{ \|\psi[\alpha]\|_\infty^2 \cdot |R(\psi[\alpha])| \} \leq \underline{B}[\alpha].$$

Indeed, let $\psi[\alpha] = \psi[\alpha]_{\ell,k,j}$ be associated with the folder $X[\alpha]_k^\ell$. As in (9.3), we have $X[\alpha]_k^\ell = \biguplus_{j \in \text{subfolders}[\alpha](\ell,k)} X[\alpha]_j^{\ell+1}$. By definition, $\psi[\alpha]$ is constant on each set $X[\alpha]_j^{\ell+1}$, with value we denote by $\psi[\alpha](X[\alpha]_j^{\ell+1})$. Let

$$j^* = \operatorname{argmax}_{j \in \text{subfolders}[\alpha](\ell,k)} \left| \psi[\alpha]_{\ell,k,j} \left(X[\alpha]_j^{\ell+1} \right) \right|,$$

so that $\|\psi[\alpha]\|_\infty = \psi[\alpha](X[\alpha]_{j^*}^{\ell+1})$. Since

$$1 = \int_X |\psi[\alpha]_{\ell,k,j}|^2 = \sum_{j \in \text{subfolders}[\alpha](\ell,k)} \left(\psi[\alpha](X[\alpha]_j^{\ell+1}) \right)^2 |X[\alpha]_j^{\ell+1}|$$

we have by (9.4)

$$\|\psi[\alpha]\|_\infty^2 = \psi[\alpha] \left(X[\alpha]_{j^*}^{\ell+1} \right)^2 \leq \frac{1}{|X[\alpha]_{j^*}^{\ell+1}|} \leq \frac{1}{B[\alpha] \cdot |X[\alpha]_k^\ell|}.$$

Multiplying over $\alpha = 1 \dots d$ yields 9.5.

The lemma now follows as

$$\int_X |\psi|^p \leq (\|\psi\|_\infty^2)^{\frac{p}{2}} \cdot |R(\psi)| \leq \left(\frac{\tau(\Psi)}{|R(\psi)|} \right)^{\frac{p}{2}} |R(\psi)| = \tau(\Psi)^{\frac{p}{2}} |R(\psi)|^{1-\frac{p}{2}}$$

□

9.6 Bounded ℓ_p -Entropy Implies Function Approximation and Decomposition

9.6.1 Problem Statement: Approximating a Function on a Product Space by a Few Terms in Its Tensor-Haar Expansion

Let X be a product space and Ψ a tensor-Haar-like basis as in Sect. 9.5.2. We are interested in conditions on a function $f : X \rightarrow \mathbb{R}$ that will allow approximating in using a small number terms in its tensor-Haar expansion, $\sum_{i=1}^N \langle f, \psi_i \rangle \psi_i$.

In [20], J. O. Strömberg, addressing the continuous version of this question, observed it is very useful to sort the tensor-Haar functions according to $|R|$, the volume of their support. He proved that if $f \in L_1([0, 1]^d)$ has bounded mixed derivatives, then

$$\sup_{x \in [0, 1]^d} \left| f(x) - \sum_{R \in \mathcal{R} \text{ s.t. } |R| > \varepsilon} \langle f, \psi_R \rangle \psi_R(x) \right| < \text{const} \cdot \varepsilon \cdot \log^{d-1} \left(\frac{1}{\varepsilon} \right),$$

where ψ_R is the unique classical tensor-Haar basis function supported on the dyadic rectangle R (see Sect. 9.6.5 below for the details of the continuous Euclidean version). However, the mixed derivative condition has certain disadvantages as a measure of function regularity: (1) It is not invariant to rotations and other simple coordinate changes, and (2) it does not generalize to general datasets. We are interested in weaker conditions, which would lead to data analysis algorithms.

We now show that one such condition is that the quantity $\sum_{i=1}^N |\langle f, \psi_i \rangle|^p$ (for some $0 < p < 2$), which we call the “ ℓ_p entropy of the tensor-Haar coefficients”, is small. This condition is easy to check and appears to be much better adapted to our setting than the mixed derivative condition.

9.6.2 An Approximation Theorem

All the integrals on X are with respect to the normalized counting measure $|\cdot|$.

Theorem 9.1. Fix $0 < p < 2$ and $f : X \rightarrow \mathbb{R}$.

$$e_p(f) = \sum_{i=1}^N |\langle f, \psi_i \rangle|^p.$$

Let $\varepsilon > 0$ and denote by $A_\varepsilon f$ an approximation of f obtained by retaining only coefficients of tensor-Haar functions, which are (i) large, and (ii) correspond to basis functions supported on large folders:

$$A_\varepsilon f = \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| > \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle f, \psi_i \rangle \psi_i(x).$$

Then

1. The number of coefficients retained in $A_\varepsilon f$ does not exceed $\varepsilon^{-1} e_p(f)$. In particular, it depends on the dimension d only through $e_p(f)$.
2. Approximation in the mean when $0 < p \leq 1$: if $e_p(f) \leq 1$ then

$$\left(\int_X |A_\varepsilon f - f|^p \right)^{\frac{1}{p}} \leq \left(\tau(\Psi)^{\frac{p}{2}} + 1 \right)^{\frac{1}{p}} \cdot \varepsilon^{\left(\frac{1}{p} - \frac{1}{2}\right)}. \quad (9.6)$$

3. Approximation in L_p when $1 \leq p < 2$:

$$\left(\int_X |A_\varepsilon f - f|^p \right)^{\frac{1}{p}} \leq \left(\tau(\Psi)^{\frac{1}{p} - \frac{1}{2}} + 1 \right) \cdot \varepsilon^{\left(\frac{1}{p} - \frac{1}{2}\right)} \cdot (e_p(f))^{\frac{1}{p}}, \quad (9.7)$$

where $(e_p(f))^{\frac{1}{p}}$ is the ℓ_p norm of the coefficient vector $\{\langle f, \psi_i \rangle\}_{i=1}^N$.

4. Uniform pointwise approximation on a set of large measure: For any $\lambda > 0$ and $1 \leq p < 2$ we have

$$|A_\varepsilon f(x) - f(x)| < \lambda \cdot \varepsilon^{\frac{1}{p} - \frac{1}{2}}$$

for any x outside an exceptional set E_λ with

$$|E_\lambda| \leq \frac{1}{\lambda^p} \left(\tau(\Psi)^{\frac{1}{p} - \frac{1}{2}} + 1 \right)^p \cdot e_p(f).$$

Proof. For part (1), the number of coefficients, we have

$$\# \left\{ 1 \leq i \leq N \mid |\langle f, \psi_i \rangle| > \varepsilon^{\frac{1}{p}} \right\} \cdot \varepsilon \leq \sum_{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| > \varepsilon^{\frac{1}{p}}} |\langle f, \psi_i \rangle|^p \leq e_p(f)$$

and hence

$$\# \left\{ 1 \leq i \leq N \mid |\langle f, \psi_i \rangle| > \varepsilon \right\} \leq \varepsilon^{-1} \cdot e_p(f).$$

The rest of our proof relies on the following key inequalities.

First key inequality: For $0 < p < 2$ we have

$$\int_X \left| \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle f, \psi_i \rangle \psi_i \right|^p \leq \varepsilon^{1-\frac{p}{2}} \cdot (e_p(f))^{\frac{p}{2}} \quad (9.8)$$

Second key inequality: for $0 < p \leq 1$ we have

$$\int_X \left| \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} \langle f, \psi_i \rangle \psi_i \right|^p \leq \tau(\Psi)^{\frac{p}{2}} \cdot \varepsilon^{1-\frac{p}{2}} \cdot e_p(f) \quad (9.9)$$

Third key inequality (3): for $1 \leq p < 2$ we have

$$\int_X \left| \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} \langle f, \psi_i \rangle \psi_i \right|^p \leq (\tau(\Psi) \cdot \varepsilon)^{1-\frac{p}{2}} \cdot e_p(f) \quad (9.10)$$

Let us first deduce the theorem from these inequalities.

For part (2), approximation in the mean, assume $0 < p \leq 1$. Recall that $|x + y|^p \leq |x|^p + |y|^p$. We have

$$\begin{aligned} \int_X |A_\varepsilon f - f|^p &= \int_X \left| f - \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| > \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle f, \psi_i \rangle \psi_i \right|^p \\ &= \int_X \left| \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{or } |R(\psi_i)| \leq \varepsilon}} \langle f, \psi_i \rangle \psi_i \right|^p \\ &= \int_X \left| \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} \langle f, \psi_i \rangle \psi_i + \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle f, \psi_i \rangle \psi_i \right|^p \\ &\leq \int_X \left| \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} \langle f, \psi_i \rangle \psi_i \right|^p + \int_X \left| \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle f, \psi_i \rangle \psi_i \right|^p. \end{aligned}$$

As $e_p(f) \leq 1$, it now follows from (9.9) and (9.8) that

$$\begin{aligned} \int_X |A_\varepsilon f - f|^p &\leq \tau(\Psi)^{\frac{p}{2}} \cdot \varepsilon^{1-\frac{p}{2}} \sum_{\substack{1 \leq i \leq N \text{ s.t.} \\ |R(\psi_i)| \leq \varepsilon}} |\langle f, \psi_i \rangle|^p + \varepsilon^{1-\frac{p}{2}} \left(\sum_{\substack{1 \leq i \leq N \text{ s.t.} \\ |R(\psi_i)| > \varepsilon}} |\langle f, \psi_i \rangle|^p \right)^{\frac{p}{2}} \\ &\leq \left(\tau(\Psi)^{\frac{p}{2}} + 1 \right) \cdot \varepsilon^{(1-\frac{p}{2})} \end{aligned}$$

which is equivalent to (9.6).

Turning to part (3), approximation in L_p , assume $1 \leq p < 2$ and let $f : X \rightarrow \mathbb{R}$ such that $e_p(f) \leq 1$. Consider the function space $L_p = L_p(X)$ (w.r.t the normalized product counting measure $|\cdot|$) and let ℓ_p denote the sequence space \mathbb{R}^N with the norm $\|(a_1 \dots a_N)\|_{\ell_p} = (\sum_{i=1}^N |a_i|^p)^{\frac{1}{p}}$. Denote by $\|f\|_{L_p} = (\int_X |f|^p)^{\frac{1}{p}}$ the norm of $L_p(X)$. Then by the inequalities (9.8) and (9.10),

$$\begin{aligned} \|A_\varepsilon f - f\|_{L_p} &\leq \left\| \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} \langle f, \psi_i \rangle \psi_i \right\|_{L_p} + \left\| \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle f, \psi_i \rangle \psi_i \right\|_{L_p} \\ &\leq (\tau(\Psi) \cdot \varepsilon)^{\frac{1}{p}-\frac{1}{2}} + \varepsilon^{\frac{1}{p}-\frac{1}{2}}. \end{aligned}$$

It follows that the norm of linear operator $\ell_p \rightarrow L_p$ defined by $(\langle f, \psi_1 \rangle, \dots, \langle f, \psi_N \rangle) \mapsto (A_\varepsilon f - f)$ is bounded by $(\tau(\Psi) + 1)^{\frac{1}{p}-\frac{1}{2}} \cdot \varepsilon^{\frac{1}{p}-\frac{1}{2}}$, as required. Finally, for part (4), uniform pointwise approximation on a set of large measure, let $\lambda > 0$ and $0 < p < 2$. Define

$$E_{\lambda,p} = \left\{ x \in X \mid |A_\varepsilon f(x) - f(x)| \geq \lambda \varepsilon^{\frac{1}{p}-\frac{1}{2}} \right\}.$$

By Markov's inequality and part (3) above we have

$$\begin{aligned} |E_{\lambda,p}| &\leq \frac{\int_X |A_\varepsilon f - f|^p}{\lambda^p \cdot \varepsilon^{1-\frac{p}{2}}} \\ &\leq \frac{1}{\lambda^p} \left(\tau(\Psi)^{\frac{1}{p}-\frac{1}{2}} + 1 \right)^p \cdot e_p(f). \end{aligned}$$

To complete the proof, we turn to the three key inequalities (9.8), (9.9) and (9.10).

To see the first key inequality (9.8), by Parseval's identity we have

$$\begin{aligned}
 \int_X \left| \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle f, \psi_i \rangle \psi_i \right|^2 &= \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} |\langle f, \psi_i \rangle|^2 \\
 &\leq \left(\varepsilon^{\frac{1}{p} - \frac{1}{2}} \right)^2 \cdot \sum_{\substack{1 \leq i \leq N \text{ s.t. } |\langle f, \psi_i \rangle| \leq \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} |\langle f, \psi_i \rangle|^p \\
 &\leq \left(\varepsilon^{\frac{1}{p} - \frac{1}{2}} \right)^2 \cdot \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| > \varepsilon} |\langle f, \psi_i \rangle|^p.
 \end{aligned}$$

Now, Hölder's inequality implies that $(\int_X |g|^p)^{\frac{1}{p}} \leq \left(\int_X |g|^2 \right)^{\frac{1}{2}}$, and (9.8) follows.

To see the second key inequality (9.9), assume that $0 < p \leq 1$. Since $|x + y|^p \leq |x|^p + |y|^p$, by Lemma 9.1 we have

$$\begin{aligned}
 \int_X \left| \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} \langle f, \psi_i \rangle \psi_i \right|^p &\leq \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} |\langle f, \psi_i \rangle|^p \int_X |\psi_i|^p \\
 &\leq \tau(\Psi)^{\frac{p}{2}} \cdot \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} |\langle f, \psi_i \rangle|^p |R(\psi_i)|^{1 - \frac{p}{2}} \\
 &\leq \tau(\Psi)^{\frac{p}{2}} \cdot \varepsilon^{1 - \frac{p}{2}} \cdot \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} |\langle f, \psi_i \rangle|^p \quad (9.11)
 \end{aligned}$$

which is (9.9). Finally, to see the third key inequality (9.10), let $1 \leq p < 2$. Consider the function space $L_p = L_p(X)$ and the sequence space ℓ_p as in the proof of part (3) above. For a sequence $\mathbf{a} = (a_1 \dots a_N)$, define

$$T(\mathbf{a}) = \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} a_i \psi_i.$$

For $0 < p \leq 2$, let $\|T\|_{p,p}$ denote the operator norm of the linear operator $T : \ell_p \rightarrow L_p$. Clearly, $\|T\|_{2,2} \leq 1$. By the inequality (9.9), we have $\|T\|_{1,1} \leq (\tau(\Psi) \cdot \varepsilon)^{\frac{1}{2}}$. We now appeal to the Riesz–Thorin Interpolation Theorem (see e.g. [19] pp. 179, Theorem 1.3), whereby for any $1 \leq p < 2$,

$$\|T\|_{p,p} \leq (\tau(\Psi) \cdot \varepsilon)^{\frac{1}{2}(1-t)},$$

whenever $0 \leq t \leq 1$ satisfies $\frac{1}{p} = 1 - t + \frac{t}{2}$, or equivalently $1 - t = \frac{2}{p} - 1$. It follows that $\|T\|_{p,p} \leq (\tau(\Psi) \cdot \varepsilon)^{\frac{1}{p} - \frac{1}{2}}$. Let $1 \leq p < 2$. We have

$$\begin{aligned} \left(\int_X \left| \sum_{1 \leq i \leq N \text{ s.t. } |R(\psi_i)| \leq \varepsilon} \langle f, \psi_i \rangle \psi_i \right|^p \right)^{\frac{1}{p}} &= \left\| T \left(\langle f, \psi_1 \rangle, \dots, \langle f, \psi_N \rangle \right) \right\|_{L_p} \\ &\leq \|T\|_{p,p} \left\| (\langle f, \psi_1 \rangle, \dots, \langle f, \psi_N \rangle) \right\|_{\ell_p} \\ &= (\tau(\Psi) \cdot \varepsilon)^{\frac{1}{p} - \frac{1}{2}} \cdot (e_p(f))^{\frac{1}{p}}. \end{aligned}$$

which is equivalent to (9.10). \square

In order to find the coefficients retained in the approximation above computationally, namely to locate $i \in \{1 \dots N\}$ such that $|\langle f, \psi_i \rangle| > \varepsilon^{\frac{1}{p}}$ and $R(\psi_i) > \varepsilon$, we must check the size of coefficients $\langle f, \psi_i \rangle$ such that $R(\psi_i) > \varepsilon$. In order to bound the number of coefficients whose size must be examined, we now show that

$$\#\{1 \leq i \leq N \mid R(\psi_i) > \varepsilon\} \leq (\beta - 1)^d \beta^2 \cdot \left(\frac{1}{\varepsilon}\right) \cdot \left(\log_{\beta} \left(\frac{1}{\varepsilon}\right) + 1\right)^{d-1}$$

where $\mathcal{T}[i]$ is $(\underline{B}[i], \overline{B}[i])$ -balanced as in (9.4), and $\beta = \max_{1 \leq i \leq d} \left\{ \frac{1}{\underline{B}[i]} \right\}$. Indeed, by (9.4), for any $(\underline{B}, \overline{B})$ partition tree, $n(\ell)$ (the number of folders in level ℓ) satisfies $n(\ell) \leq \left(\frac{1}{\underline{B}}\right)^{\ell-1}$. Denote by $(\underline{B}[i], \overline{B}[i])$ the parameters for the partition tree $\mathcal{T}[i]$ and let $\beta = \max_{1 \leq i \leq d} \left\{ \frac{1}{\underline{B}[i]} \right\}$. If $\beta^{-L} \leq \varepsilon < \beta^{-L+1}$, then

$$\#\{R \in \mathcal{R} \text{ s.t. } |R| \geq \beta^{-L}\} = \sum_{\substack{(r_1, \dots, r_d) \in \mathbb{N}^d \\ \beta^{-\sum r_i} \geq \beta^{-L}}} \beta^{\sum r_i} = \sum_{\substack{(r_1, \dots, r_d) \in \mathbb{N}^d \\ r_1 + \dots + r_d \leq L}} \beta^{\sum r_i}.$$

Let us show by induction that $\sum_{\substack{(r_1, \dots, r_d) \in \mathbb{N}^d \\ r_1 + \dots + r_d \leq L}} \beta^{\sum r_i} \leq \beta^{L+1} L^{d-1}$. Indeed, for $d=1$ we have $\sum_{r=0}^L \beta^r = \beta^{L+1} - 1$. Assuming this for d , we have

$$\sum_{\substack{(r_1, \dots, r_{d+1}) \in \mathbb{N}^{d+1} \\ r_1 + \dots + r_{d+1} \leq L}} \beta^{(r_1 + \dots + r_{d+1})} = \sum_{r_{d+1}=0}^L \beta^{r_{d+1}} \sum_{\substack{(r_1, \dots, r_d) \in \mathbb{N}^d \\ r_1 + \dots + r_d \leq L - r_{d+1}}} \beta^{(r_1 + \dots + r_d)}$$

$$\begin{aligned}
&\leq \sum_{r_{d+1}=0}^L \beta^{r_{d+1}} \left(\beta^{L-r_{d+1}+1} \cdot (L-r_{d+1})^{d-1} \right) \\
&= \beta^{L+1} \sum_{k=0}^{L-1} k^{d-1} \leq \beta^{L+1} L \cdot L^{d-1}.
\end{aligned}$$

Therefore, $\#\{R \in \mathcal{R} \text{ s.t. } |R| > \varepsilon\} \leq \beta \cdot \beta^L \cdot L^{d-1} \leq \beta^2 \cdot \left(\frac{1}{\varepsilon}\right) \cdot \left(\log_{\beta} \left(\frac{1}{\varepsilon}\right) + 1\right)^{d-1}$. Since up to $(\beta - 1)^d$ can be associated with any rectangle $R \in \mathcal{R}$, we finally obtain that

$$\#\{1 \leq i \leq N \mid |R(\psi_i)| > \varepsilon\} \leq (\beta - 1)^d \beta^2 \cdot \left(\frac{1}{\varepsilon}\right) \cdot \left(\log_{\beta} \left(\frac{1}{\varepsilon}\right) + 1\right)^{d-1}.$$

A tighter bound can be obtained by considering each $\underline{B}[i]$ separately instead of using the uniform bound β .

9.6.3 ℓ_1 -Entropy Interpreted as Smoothness

Consider a data matrix and suppose we have constructed a two coupled graphs, one on the rows and one of the columns. Did we do a good job? Theorem 9.1 implies that the ℓ_p entropy of the matrix in the tensor-Haar-like basis, induced by two partition trees describing the two geometries, provides a way to quantify the usefulness of the coupled geometry at hand. In general, we would like to be able to say that a function is adapted to a given geometry if the function is smooth and simple in this geometry. Instead of smoothness in the Euclidean sense of differentiability, in our general setting it becomes natural to quantify *smoothness* using a Haar-like basis in terms of pointwise exponential decay of coefficients (see [12] for a theorem of this kind), while *simplicity* means a small number of non-negligible tensor-Haar-like coefficients, which is, in some cases, intimately related to small ℓ_1 -entropy. The above inequalities for Haar-like bases relate function smoothness in the geometry, as it is captured by the basis, to sparsity of the Haar expansion. The unique feature of the tensor-Haar-like basis is that *only coefficients of Haar functions ψ_i with support volume $|R(\psi_i)| > \varepsilon$ are required*, eliminating the need to consider coefficients that correspond to basis functions with small support.

In particular, we find an answer to Question 1 from Sect. 9.2.2.

9.6.4 A Decomposition Theorem

The relationship between coefficient decay and smoothness is well-known in harmonic analysis. In [12, 13] it is shown that, for Haar-like bases, coefficient decay

is equivalent to Hölder property, just as in the classical case. We now extend this fact to Tensor-Haar-like bases and deduce a decomposition theorem of Calderón–Zygmund type.

For a set $X[i]$ with a partition tree $\mathcal{T}[i]$, we define a metric on $X[i]$ by

$$\rho(x, y) = \begin{cases} |folder(x, y)| & x \neq y \\ 0 & x = y \end{cases} \quad (9.12)$$

where $folder(x, y)$ is the smallest folder in any level of $\mathcal{T}[i]$ containing both x and y .

Claim. Let $f : X[i] \rightarrow \mathbb{R}$ and $p > 0$. Let Ψ be a Haar-like basis corresponding to the partition tree $\mathcal{T}[i]$. Consider the conditions:

1. There is a constant C such that for any $\psi \in \Psi$ associated with the folder X_k^ℓ ,

$$|\langle f, \psi \rangle| \leq C \cdot |X_k^\ell|^{p+1/2}.$$

2. f is p -Hölder with constant C in the metric ρ . That is, $|f(x) - f(y)| \leq C \cdot \rho(x, y)^p$ for all $x, y \in X$.

Then (1) with constant $C \implies$ (2) with constant γC , and (2) with constant $C \implies$ (1) with constant δC . The numbers γ and δ depend on p and the partition tree \mathcal{T} alone.

See [13] for the proof.

In our setting of a product space, a similar phenomenon occurs. We discuss the case of matrices, where $d = 2$.

Theorem 9.2. $X = X[1] \times X[2]$. Suppose that $\mathcal{T}[i]$ is a partition tree on $X[i]$ and Ψ_i is an induced Haar-like basis, $i = 1, 2$. Let Ψ be the tensor-Haar-like basis. Let $f : X \rightarrow \mathbb{R}$ and $p > 0$. Consider the following conditions:

1. There is a constant C such that for any $\psi \in \Psi$ associated with the rectangle $R(\psi)$,

$$|\langle f, \psi \rangle| \leq C \cdot |R(\psi)|^{p+1/2}.$$

2. f is p -bi-Hölder with constant C in the metric ρ , in the sense that

$$|f(x, y) - f(x', y) - f(x, y') + f(x', y')| \leq C \cdot \rho[1](x, x')^p \rho[2](y, y')^p$$

for some constant C , all $x, x' \in X[1]$ and all $y, y' \in X[2]$. Here, $\rho[i]$ is the metric induced by $\mathcal{T}[i]$ on $X[i]$, as in (9.12), for $i = 1, 2$.

Then (1) with constant $C \implies$ (2) with constant γC , and (2) with constant $C \implies$ (1) with constant δC , where the numbers γ, δ depends on p and on the partition trees $\mathcal{T}[1]$ and $\mathcal{T}[2]$ alone.

Proof. \implies Assume (1) holds with constant C . Choose $x, x' \in X[1]$ and $y, y' \in X[2]$. Write the tensor-Haar-like function explicitly as products: $\Psi = \{\psi_i(x)\varphi_j(y)\}_{i,j}$ for $1 \leq i \leq m$ and $1 \leq j \leq n$. Then

$$f(x, y) = \sum_{i,j} \langle f, \psi_i \varphi_j \rangle \psi_i(x) \varphi_j(y) = \sum_i a_i(y) \psi_i(x),$$

where $a_i(y) = \sum_j \langle f, \psi_i \varphi_j \rangle \varphi_j(y)$. Since $|\langle f, \psi_i \varphi_j \rangle| \leq C \cdot |R(\psi_i)|^{p+1/2} |R(\varphi_j)|^{p+1/2}$, where $R(\varphi_j)$ is the folder in $\mathcal{T}[2]$ associated with φ_j , by Theorem 9.2 above for $X[2]$, we have

$$|a_i(y) - a_i(y')| \leq \gamma_2 \cdot C \cdot |R(\psi_i)|^{p+1/2} \cdot \rho[2](y, y')^p,$$

where γ_2 depends on p and $\mathcal{T}[2]$ alone. Now, define $g(x) = f(x, y) - f(x, y') = \sum_i [a_i(y) - a_i(y')] \psi_i(x)$. Again by Theorem 9.2, this time for $X[1]$,

$$|g(x) - g(x')| \leq \gamma_1 \cdot \gamma_2 \cdot C \cdot \rho[1](x, x')^p \cdot \rho[2](y, y')^p.$$

\Leftarrow Assume (2) holds with constant C . As $\int_{X_1} \psi_i(x) = 0$ and similarly for X_2 , we have for any $\psi = \psi_i(x)\varphi_j(y)$ and any $(x', y') \in R(\psi)$ that

$$\begin{aligned} \langle f, \psi_i \varphi_j \rangle &= \int_X f(x, y) \psi_i(x) \varphi_j(y) dx dy \\ &= \int_X [f(x, y) - f(x', y) - f(x, y') + f(x', y')] \psi_i(x) \varphi_j(y) dx dy, \end{aligned}$$

hence

$$|\langle f, \psi_i \varphi_j \rangle| \leq C \cdot \rho[1](x, x')^p \rho[2](y, y')^p \cdot \int_X |\psi_i(x) \varphi_j(y)| dx dy.$$

But by the definition of $\rho[1]$, we have $\rho[1](x, x') \leq R(\psi_i)$ and similarly, $\rho[2](y, y') \leq R(\varphi_j)$. Finally, by Lemma (9.1), we have

$$\int_X |\psi_i(x) \varphi_j(y)| dx dy \leq \sqrt{\tau(\Psi)} |R(\psi_i)|^{\frac{1}{2}} |R(\varphi_j)|^{\frac{1}{2}}.$$

In summary,

$$|\langle f, \psi_i \varphi_j \rangle| \leq \sqrt{\tau(\Psi)} \cdot C \cdot |R(\psi)|^{p+\frac{1}{2}}.$$

□

We can decompose a given function on X to a regular part and a part with small support, à la Calderón-Zygmund. In the context of a data matrix or tensor, this

would mean that if we can describe f efficiently using a tensor-Haar-like basis, then it decomposes into a “typical” matrix, which is regular with respect to the coupled geometry we constructed, and an “outlier” matrix or potentially irregular behavior but with small support.

Theorem 9.3. *Let $f : X \rightarrow \mathbb{R}$ and $0 < p < 2$. There is a decreasing sequence of sets $E_\ell \subset X$ where $|E_\ell| < \frac{e_p(f)}{2^\ell}$ for any ℓ , and a decomposition $f = g_\ell + b_\ell$, such that $x \mapsto g_\ell(x)$ (the “good” function) is $\frac{1}{p} - \frac{1}{2}$ -Hölder with constant $2^{\frac{\ell}{p}}$ and $x \mapsto b_\ell(x)$ (the “bad” function) is supported on E_ℓ . The functions g_ℓ and b_ℓ , as well as the set E_ℓ , all have explicit form.*

Proof. Define

$$S_p(x) = \sum_{i=1}^N |\langle f, \psi_i \rangle|^p \frac{\chi_{R(\psi_i)}(x)}{|R(\psi_i)|}$$

where $\chi_{R(\psi_i)}(x)$ is the indicator function of $R(\psi_i)$. (This is an analog of the Littlewood–Paley function). A Fubini argument gives

$$\int_X S_p(x) dx = \sum_{i=1}^N |\langle f, \psi_i \rangle|^p = e_p(f).$$

As in the proof of Theorem 9.1, define an “exceptional set”

$$E_{\ell,p} = \left\{ x \in X \mid S_p(x) > 2^\ell \right\}.$$

Clearly $E_{\ell+1,p} \subset E_{\ell,p}$. Markov’s inequality gives

$$|E_{\ell,p}| \leq \frac{e_p(f)}{2^\ell}.$$

Now decompose the set of all rectangles \mathcal{R} as follows. Define

$$\mathcal{R}_\ell = \left\{ R \in \mathcal{R} \mid R \subseteq E_{\ell,p} \text{ and } R \not\subseteq E_{\ell+1,p} \right\}$$

and observe that $\mathcal{R} = \biguplus_{\ell \in \mathbb{Z}} \mathcal{R}_\ell$. We can, thus, decompose $f = \sum_{\ell \in \mathbb{Z}} f_\ell$, where

$$f_\ell = \sum_{1 \leq i \leq N \text{ with } R(\psi_i) \in \mathcal{R}_\ell} \langle f, \psi_i \rangle \psi_i.$$

Fix $\ell \in \mathbb{Z}$ and $R \in \mathcal{R}_\ell$. By the definition of \mathcal{R}_ℓ , there exists $x \in R$ with $x \in E_{\ell,p} \setminus E_{\ell+1,p}$. For this x we have

$$\sum_{i=1}^N |\langle f, \psi_i \rangle|^p \frac{\chi_{R(\psi_i)}(x)}{|R(\psi_i)|} = S_p(x) \leq 2^{\ell+1}.$$

Choose $1 \leq i \leq N$ with $R(\psi_i) = R$. As $\chi_{R(\psi_i)}(x) = 1$, we get

$$\frac{|\langle f, \psi_i \rangle|^p}{|R(\psi_i)|} \leq S_p(x) \leq 2^{\ell+1},$$

namely $|\langle f, \psi_i \rangle| \leq 2^{\frac{\ell+1}{p}} |R(\psi_i)|^{\frac{1}{p}}$. Define $g_\ell = \sum_{k=1}^{\ell-1} f_k$ and $b_\ell = \sum_{k \geq \ell} f_k$. By Theorem 9.2, g_ℓ is as required. Clearly, b_ℓ is supported on $E_{\ell,p}$. \square

We note that the above is not an existence result, but rather an *explicit* formula for decomposing a given function.

9.6.5 The Euclidean Analog

We briefly translate the above general, discrete results to the Euclidean case of $f \in L_1([0, 1]^d)$ and the tensor product of classical Haar bases. This yields an approximation theorem in high dimensional Euclidean spaces.

Consider the unit interval $[0, 1] \subset \mathbb{R}$ and the classical Haar basis in $[0, 1]$. Each Haar functions corresponds to its support, a dyadic interval. We index the basis functions by $\{\psi_I\}_{I \in \mathcal{I}}$ where $\mathcal{I} = \{[n2^{-k}, (n+1)2^{-k}]\}_{n,k}$. Consider the unit cube $[0, 1]^d$ in \mathbb{R}^d . The *tensor-Haar basis* on $[0, 1]^d$, namely $\{\psi_{I_1} \otimes \dots \otimes \psi_{I_d}\}_{I_i \in \mathcal{I}}$, is an orthonormal basis for $L_2([0, 1]^d)$. Let \mathcal{R} denote the set of dyadic rectangles on $[0, 1]^d$, $\mathcal{R} = \{I_1 \times \dots \times I_d \mid I_i \in \mathcal{I}\}$. As each function in the tensor-Haar basis corresponds to its support, a dyadic rectangle, we index the tensor-Haar basis by $\{\psi_R\}_{R \in \mathcal{R}}$, where for a dyadic rectangle $R = I_1 \times \dots \times I_d$, we have $\psi_R = \psi_{I_1} \otimes \dots \otimes \psi_{I_d}$. Write $|R|$ for the d -dimensional Lebesgue measure of R .

The proof of Theorem 9.1 can be used verbatim to prove its Euclidean version:

Theorem 9.4. *For $f \in L_1([0, 1]^d)$ and $0 < p < 2$, write*

$$e_p(f) = \sum_{R \in \mathcal{R}} |\langle f, \psi_R \rangle|^p.$$

Let $\varepsilon > 0$ and denote by Af an approximation of f obtained by retaining only large coefficients of tensor-Haar functions supported on large folders:

$$A_\varepsilon f = \sum_{\substack{R \in \mathcal{R} \text{ s.t. } |\langle f, \psi_R \rangle| > \varepsilon^{\frac{1}{p}} \\ \text{and } |R| > \varepsilon}} \langle f, \psi_R \rangle \psi_R(x).$$

Then

1. The number of coefficients retained in $A_\varepsilon f$ does not exceed $\varepsilon^{-1} e_p(f)$.
2. Approximation in the mean when $0 < p \leq 1$: if $e_p(f) \leq 1$ then

$$\left(\int_{[0,1]^d} |A_\varepsilon f - f|^p \right)^{\frac{1}{p}} \leq 2^{\frac{1}{p}} \cdot \varepsilon^{\left(\frac{1}{p}-\frac{1}{2}\right)}. \quad (9.13)$$

3. Approximation in L_p when $1 \leq p < 2$:

$$\left(\int_{[0,1]^d} |A_\varepsilon f - f|^p \right)^{\frac{1}{p}} \leq 2 \cdot \varepsilon^{\left(\frac{1}{p}-\frac{1}{2}\right)} \cdot (e_p(f))^{\frac{1}{p}}. \quad (9.14)$$

4. Uniform pointwise approximation on a set of large measure: For any $\lambda > 0$ and $1 \leq p < 2$ we have

$$|A_\varepsilon f(x) - f(x)| < \lambda \cdot \varepsilon^{\frac{1}{p}-\frac{1}{2}}$$

for any x outside an exceptional set E_λ with

$$|E_\lambda| \leq \frac{2^p}{\lambda^p} e_p(f).$$

5. We only need to evaluate the size of

$$\#\{R \in \mathcal{R} \mid |R| > \varepsilon\} \leq 4 \cdot \left(\frac{1}{\varepsilon}\right) \cdot \left(\log\left(\frac{1}{\varepsilon}\right) + 1\right)^{d-1}$$

coefficients.

We remark that the class of functions, for which the Haar expansion has finite entropy has been characterized by [8]¹.

The relation between smoothness and tensor-Haar coefficient size of Theorem 9.2 is well known in the Euclidean case. Theorem 9.3 extends to the Euclidean case without change in its statement and proof.

¹ It is shown there that this class is independent of the wavelet used, and is equivalent to having a Harmonic extension whose derivative is integrable in the disk (or bi-disk). They also characterize the dual spaces as Bloch spaces, which in our case are just functions with bounded Haar coefficients. Observe further that

$$f = \sum_R |R|^{\frac{1}{2}} \langle f, \psi_R \rangle |R|^{-\frac{1}{2}} \psi_R(x)$$

is a special atomic decomposition of $\left| \frac{d}{dx} \right|^{\frac{1}{2}} \left| \frac{d}{dy} \right|^{\frac{1}{2}} f$, which is therefore in the Hardy space H^1 of the bi-disk. A similar result holds for other entropies, implying a fractional derivative in the Hardy space.

9.7 Database Analysis

The iterative procedure for recovering the coupled geometry of a data matrix, together with the approximation theorem, suggest a family of data analysis schemes. The general recipe is as follows.

Algorithm 9.2. Given a data matrix M ,

1. Compute an initial affinity W_0 on the columns of M .
2. Apply an iteration of Algorithm 9.1 and obtain partition trees $\mathcal{T}_{\text{rows}}$ and $\mathcal{T}_{\text{cols}}$ on the rows and columns of M , respectively.
3. Construct induced Haar-like bases Ψ_{rows} and Ψ_{cols} and the tensor-Haar-like basis $\Psi = \{\psi_i\}_{i=1}^N$.
4. Compute the ℓ_1 entropy $e_1(M, \Psi) = \sum_{i=1}^N |\langle M, \psi_i \rangle|$.
5. Repeat steps 2–4 until $e_1(M, \Psi)$ converges.
6. Transform² M into the tensor-Haar-like basis Ψ to obtain the coefficient matrix \tilde{M} .
7. Process M in the coefficient domain \tilde{M} (see below) and transform back.

This data driven geometric “self-organization” allows to analyze any data matrix according to its intrinsic row or column structure. While this presentation focuses on data matrices, there is nothing special about order-2 tensors: this approach generalizes to data tensors of order d .

The specialization of this general procedure depends on the data analysis task at hand. Each application calls for a detailed treatment. This is beyond the scope of this introduction.

1. *Compression*: To compress M , store both partition trees $\mathcal{T}_{\text{rows}}$ and $\mathcal{T}_{\text{cols}}$, and only coefficients of tensor-Haar-like functions with $R(\psi_i) > \text{const} \cdot \varepsilon^2$, where ε is the desired error.
2. *Denoising*: Sort coefficients $\{\langle M, \psi_i \rangle\}_{i=1}^N$ according to $R(\psi_i)$ and apply thresholding.
3. *Matrix completion and missing value imputation*: If we are given the trees $\mathcal{T}_{\text{rows}}$ and $\mathcal{T}_{\text{cols}}$ from an external source or prior knowledge, or if the number of missing values in M allows reliable construction of $\mathcal{T}_{\text{rows}}$ and $\mathcal{T}_{\text{cols}}$, we can estimate the coefficients $\langle M, \psi_i \rangle$ using available entries. Using a tensor-Haar-like basis has the distinct advantage that we need only care about averages of the available points on large sets, leading to small estimator variance. See [12] for a detailed analysis in the case $d = 1$.
4. *Anomaly detection*: The residual matrix

$$\sum_{1 \leq i \leq N \text{ with } |R(\psi_i)| \leq \varepsilon} \langle M, \psi_i \rangle \psi_i$$

² Note that a fast algorithm for computing the coefficients $\langle M, \psi_i \rangle$, or equivalently, transforming M into \tilde{M} exists. We do not pursue this here.

contains the deviations from the average matrix. This yields a method for detecting anomalous observations, as they are different from the “average matrix” in their respective folders.

9.7.1 *Eliminating Haar Artifacts*

While Haar-like bases are simple to construct and to use in our general setting, their lack of smoothness introduce artifacts into any computation that is taking place in the coefficient domain. This phenomenon was studied in [5] in signal processing setting of classical Haar basis on $[0, L]$. Indeed, the location of the Haar discontinuities, namely the alignment of the dyadic grid on $[0, L]$, is arbitrary. When denoising a function $f : [0, L] \rightarrow \mathbb{R}$ using coefficient shrinkage in a given Haar basis, artifacts appear along this arbitrarily shifted dyadic grid. Therefore, they suggest to denoise f using several differently shifted Haar bases, and average the resulting denoised versions of f in order to eliminate the artifacts.

This observation is ideal for eliminating artifacts in our setting. In the general data-analysis recipe Algorithm 9.2, after completing step 5, namely after stabilizing the coupled geometry of the data matrix M , we have the affinities W_{rows} and W_{cols} on the rows and columns of M , respectively. Recall that the procedure for integrating an affinity into a partition tree, described in Sect. 9.2.1, is random. In order to reduce the artifacts caused by processing M in a tensor-Haar-like basis, we choose r and construct partition trees $\left\{ \mathcal{T}_{\text{rows}}^{(i)} \right\}_{i=1}^r$ using W_{rows} and $\left\{ \mathcal{T}_{\text{cols}}^{(i)} \right\}_{i=1}^r$ using W_{cols} . Computing induced Haar-like bases $\left\{ \Psi_{\text{rows}}^{(i)} \right\}_{i=1}^r$ and $\left\{ \Psi_{\text{cols}}^{(j)} \right\}_{j=1}^r$, this yields r^2 different Tensor-Haar-like bases $\left\{ \Psi^{(i,j)} \right\}_{i,j=1}^r$ by setting $\Psi^{(i,j)} = \Psi_{\text{rows}}^{(i)} \otimes \Psi_{\text{cols}}^{(j)}$. Each basis is used to produce a processed version of M (as in the last step of the general recipe Algorithm 9.2), and the r^2 versions are averaged to produce the final result.

We remark that one can modify this basic construction of a hierarchical scale decomposition, in order to build wavelets that provide filters restricting the frequency content of a function to bands of eigenfunctions of the diffusion or Laplace operator on the graph. See, for example, the constructions in [9, 10].

9.7.2 *Example: The Science News Database*

As a concrete example of data analysis using the general recipe Algorithm 9.2, we consider a term-document data matrix. The documents are abstracts of 1,047 articles obtained from the Science News journal website, each under one of the following scientific fields: Anthropology, Astronomy, Social Sciences, Earth Sciences, Biology, Mathematics, Medicine, or Physics. This data set was prepared and pre-processed for [17], where additional information can be found, and kindly provided by J. Solka. In the preprocessing step, a dictionary of 10,906 words were chosen

as relevant for this body of documents. Of these, the 1,000 words with the highest correlation to article subject classification were selected. The data matrix is thus $(M_{i,j})_{i=1\dots 1,000; j=1\dots 1,024}$ where the entry $M_{i,j}$ is the relative frequency of the i th word in the j th document.

The original data matrix makes little sense to the naked eye and is not shown. When running steps 1–5 in Algorithm 9.2 we observe decrease and convergence of the ℓ_1 entropy $e_1(M, \Psi) = \sum_{i=1}^N |\langle M, \psi_i \rangle|$ (Fig. 9.7). In Fig. 9.8 (left) the data




Fig. 9.7 Iterative procedure: ℓ_1 entropy over iteration number




Fig. 9.8 Science News matrix, rows and columns reordered by trees. *Left*: Original data matrix. *Right*: Approximated matrix using 0.15 of the tensor-Haar-like coefficients




Fig. 9.9 Science News ℓ_1 approximation error, theoretical bound and fraction of coefficients retained over ϵ

matrix is shown with rows and columns re-organized using depth-first order of the obtained trees. Figure 9.8 compares the original matrix with an approximation obtained by retaining those 15% of the tensor-Haar-like coefficients corresponding to $\epsilon = 9.5 \cdot 10^{-5}$. We remark that for this data matrix, as well as for other data matrices with “intrinsic organization of the row and column sets”, this approximation is superior to counterparts from classical multivariate analysis, such as retaining the 15% largest singular values in the SVD decomposition of the data matrix). Figure 9.9 shows the approximation quality $\int_X |A_\epsilon - A|$ as function of ϵ , together with the fraction of coefficients retained for the approximation and the theoretical bound of Theorem 9.1. Evidently, our bound is pessimistic – the approximation obtained is much better.

9.8 Example: Numerical Compression of a Potential Operator – Continued

Operator compression is the ability to store, apply and compute functions of very large matrices. As these numerical tasks determine the limit of many scientific computations and simulations, it is a fundamental problem of numerical analysis. When the matrix at hand describes particle interaction, such as Columbic interaction in three-dimensional space, multiscale [14] and wavelet [3] methods have proved extremely useful. (See [16] for a survey). In Sect. 9.3, we saw that we can organize a potential matrix even when the geometries involved are unknown and the rows

and columns are given in a random order. We now return to this example and show that, having organized the matrix, it can be compressed efficiently in the tensor-Haar basis. Indeed, a procedure consisting of recovering the coupled geometry of an operator and compressing it in the induced tensor-Haar basis is a natural extension of the ideas of [3, 14] to the setting where one is given only the operator describing unknown interaction on unknown geometries.

We observed that the approximation theorem yields an operator approximation scheme. Suppose that $M_{i,j}$ is the matrix of an operator and suppose we can construct a tree $\mathcal{T}_{\text{rows}}$ on the rows of M and a tree $\mathcal{T}_{\text{cols}}$ on the columns of M such that the ℓ_1 entropy of $M = \sum_{i=1}^N \langle M, \psi_i \rangle \psi_i$ in the induced tensor-Haar-like basis $\Psi = \{\psi_i\}_{i=1}^N$, namely $e_1(M) = \sum_{i=1}^N |\langle M, \psi_i \rangle|$, is small. (As before, $N = \#\text{rows} \cdot \#\text{cols}$ is the number of entries in M). Denote by M_ε the approximation

$$M_\varepsilon = \sum_{\substack{1 \leq i \leq N \text{ with } |\langle f, \psi_i \rangle| > \varepsilon^{\frac{1}{p}} \\ \text{and } |R(\psi_i)| > \varepsilon}} \langle M, \psi_i \rangle \psi_i$$

for M . From the approximation theorem, Theorem 9.1, we obtain that

$$\frac{1}{N} \sum_{x,y} |M(x,y) - M_\varepsilon(x,y)| \leq \sqrt{\varepsilon} \cdot \left(\sqrt{\tau(\Psi)} + 1 \right) e_1(M),$$

where

$$e_1(M) = \sum_{i=1}^N |\langle M, \psi_i \rangle|.$$

Returning to the potential operator example from Sect. 9.3, we find that this yields an efficient compression scheme. The tensor-Haar-like coefficients matrix is shown in Fig. 9.10. Figure 9.11 shows the $\ell_2 \rightarrow \ell_2$ operator norm of the residual $M_\varepsilon - M$ for different values of ε . Figure 9.12 shows the ℓ_1 norm of the residual and the bound from Theorem 9.1. It seems that this bound is too pessimistic.

9.9 Conclusion

Data matrices, for which both the variable/feature/covariate set and the observation set have an intrinsic organization, are very common. The framework described here proposes a data-driven geometric approach for data “self-organization” and nonparametric statistical analysis, which applies to any data set given as a product structure (of order 2 or higher). The basic assumption underlying this method is the existence of intrinsic row and column structure. This framework relies on several of observations:

1. A partition tree on a single graph, or data set, leads to Haar-like multiscale analysis of function on it [12, 13].

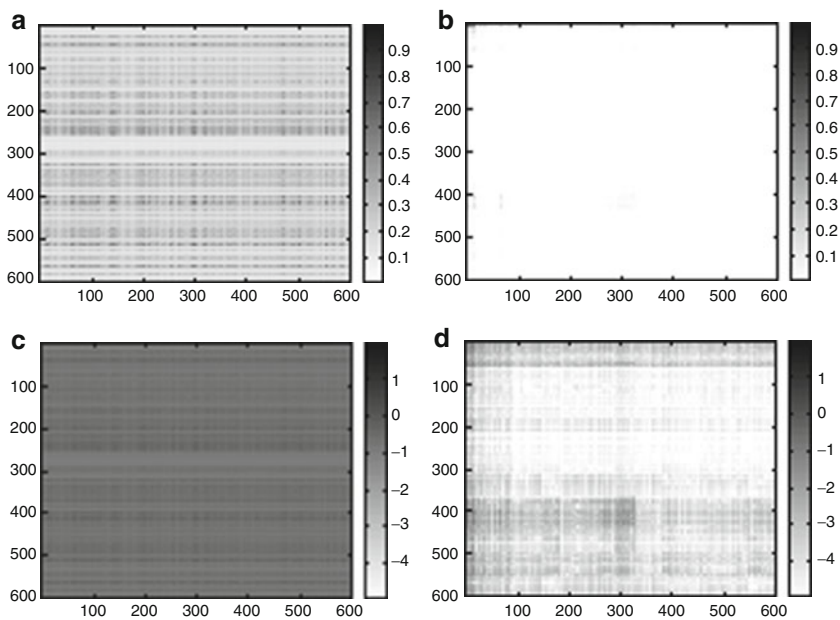


Fig. 9.10 (a) “Scrambled” potential operator matrix from Fig. 9.2. (b) Absolute value of the corresponding coefficient matrix in a tensor-Haar-like basis (color saturation set to 1 for better visibility). (c) The “Scrambled” matrix on \log_{10} scale. (d) The Absolute value of coefficient matrix on \log_{10} scale




Fig. 9.11 $\ell_2 \rightarrow \ell_2$ operator norm of the residual and number of coefficients retained over ε




Fig. 9.12 ℓ_1 norm of the error, theoretical bound and number of coefficients retained over ϵ

2. When the variables set is not arbitrary but rather carries intrinsic structure, it is useful to consider the *coupled* structure of the data matrix. Formally, we view the matrix itself as a function on the product of its row and columns spaces. This extends naturally to higher order data tensors (Sect. 9.2).
3. The first role of tensor product of Haar-like bases is that the absolute sum of expansion coefficients quantifies the regularity of the data set in a proposed coupled geometry (Sect. 9.6).
4. Their second role is for analyzing the data set. By expanding the data matrix in a “well-adapted” tensor-Haar-like basis, we can analyze the data in the coefficient domain. Tensor product of Haar-like basis offers a compact representation of data tensors of rank d that is, in some sense, not affected by d . Expansion coefficients in a tensor-Haar-like have a *one-dimensional* organization that relies on the support size $|R(\psi_i)|$, thus allowing direct use of classical signal processing ideas for high dimensional data sets (Sect. 9.7).

We remark that the tensor-Haar-like basis used here may be replaced by a tensor product of other constructions, such as scaling functions. In particular, two distinct bases can be used for the column set and for the row set of a matrix, including any Laplacian eigenfunctions of the row or columns affinity graph. From this perspective, the Singular Value Decomposition is a special case of this construction.

Acknowledgements We are indebted to our collaborator Boaz Nadler for many insights regarding Haar-like bases. We also thank Rob Tibshirani and Fred Warner for their helpful comments. M.G. is supported by a William R. and Sara Hart Kimball Stanford Graduate Fellowship and is grateful for the hospitality of the Yale Applied Math program during the preparation of this work.

References

1. Allen, G.I., Tibshirani, R.: Transposable regularized covariance models with an application to missing data imputation. To appear in *Annals of Applied Statistics* (2010)
2. Belkin, M., Niyogi, P., Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.* **13**, 1373–1397 (2003)
3. Beylkin, G., Coifman, R., Rokhlin, V., Wavelets in numerical analysis. In: Ruskai, M.B., Beylkin, G., Coifman, R. (eds.) *Wavelets and their applications*, pp. 181–210. Jones and Bartlett, Boston (1992)
4. Candès, E.J., Tao, T., Near-optimal signal recovery from random projections: Universal encoding strategies?. *IEEE Trans. Inform. Theory*, **52**(12), 5406–5425 (2006)
5. Coifman, R.R., Donoho D.L.: Translation invariant de-noising. In: Antoniadis, A., Oppenheim, G. (eds.) *Wavelets and Statistics*, pp. 125–150. Springer, New York (1995)
6. Coifman, R.R., Lafon, S.: Diffusion maps. *Appl. Comput. Harmon. Anal.* **21**(1), 5–30 (2006)
7. Coifman, R.R., Maggioni, M.: Diffusion wavelets. *Appl. Comput. Harmon. Anal.* **21**(1), 54–95 (2006)
8. Coifman, R.R., Rochberg, R.: Another characterization of B.M.O. *Proc. Amer. Math. Soc.* **79**, 249–254 (1980)
9. Coifman, R.R., Weiss, G.: *Analyse Harmonique Noncommutative sur Certains Espaces Homogenes*. Springer-Verlag (1971)
10. Coifman, R.R., Weiss, G.: Extensions of Hardy spaces and their use in analysis. *Bul. Of the A.M.S.*, **83**(4), 569–645 (1977)
11. Donoho, D.L.: Compressed Sensing. *IEEE Trans. Inform. Theory* **52**(4), 1289–1306 (2006)
12. Gavish, M., Nadler, B., Coifman, R.R.: Multiscale wavelets on trees, graphs and high dimensional data: Theory and applications to semi supervised learning. *Proceedings of the 27th International Conference on Machine Learning, ICML* (2010)
13. Gavish, M., Nadler, B., Coifman, R.R.: Inference by Haar-like wavelet analysis. preprint (2010)
14. Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations, *J. Comput. Phys.* **73**, 325–348 (1987)
15. Lazzeroni, L., and Owen, A., Plaid models for gene expression data. *Statistica Sinica* **12**(1), 61–86 (2002)
16. Martinsson, P., Tygert, M., Multilevel Compression of Linear Operators: Descendants of Fast Multipole Methods and Calderón-Zygmund Theory. Lecture notes, Yale University and Courant Institute (2009) Available at <http://cims.nyu.edu/~tygert/gradcourse/survey.pdf>. Cited 30 May 2010
17. Priebe, C.E., Marchette, D.J., Park, Y., Wegman, E.J., Solka, J.L., Socolinsky, D.A., Karakos, D., Church, K.W., Guglielmi, R., Coifman, R.R., Link, D., Healy, D.M., Jacobs, M.Q., Tsao, A.: Iterative denoising for cross-corpus discovery. *Proceedings of COMP-STAT 2004*, Physica-Verlag/Springer (2004)
18. Singh, A., Nowak, R., Calderbank, R.: Detecting weak but hierarchically-structured patterns in networks. *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics, AISTATS* (2010)
19. Stein, E.M., Weiss, G.: *Fourier analysis on Euclidean spaces*. Princeton University Press, Princeton (1971)
20. Strömberg, J.O.: Wavelets in higher dimensions. *Documenta Mathematica Extra Volume ICM-1998*(3), 523–532 (1998)
21. Tukey, J.W.: The future of data analysis, *Ann. Math. Statist.* **33**(1), 1–67 (1962)
22. Wallmann, D.M.: Multiscale diffusion coordinate refinement. Ph.D thesis, Yale University (2009)

Chapter 10

Some Recent Advances in Multiscale Geometric Analysis of Point Clouds

Guangliang Chen, Anna V. Little, Mauro Maggioni, and Lorenzo Rosasco

Abstract We discuss recent work based on multiscale geometric analysis for the study of large data sets that lie in high-dimensional spaces but have low-dimensional structure. We present three applications: the first one to the estimation of intrinsic dimension of sampled manifolds, the second one to the construction of multiscale dictionaries, called Geometric Wavelets, for the analysis of point clouds, and the third one to the inference of point clouds modeled as unions of multiple planes of varying dimensions.

10.1 Introduction

Data sets that arise in a variety of settings – from images and movies to web pages, customer transaction records, gene microarrays, etc. – are being collected at ever increasing speed and level of detail. The increase in the amount of data has not always been matched by our understanding of how to efficiently extract information, and to search, organize, and derive useful predictions from such data sets. The analysis of such data sets, modeled as point clouds in high-dimensional spaces, is an emerging and challenging area in applied mathematics, at the boundary with other disciplines such as computer science, engineering, signal processing, biology, and more. There are many applications, including organization of large libraries of documents, face recognition [52], semi-supervised learning [5, 73, 83], nonlinear image denoising and segmentation [80, 83], clustering [3, 72], machine learning [5, 67, 68, 73, 83, 84], processing of articulated images [44], cataloguing of galaxies [45], pattern analysis of brain potentials [66], the study of brain tumors [21], document classification and web searching, hyperspectral imaging, and many others. The analysis and estimation of geometric (intended in the widest sense) properties of the data include problems such as dimension estimation (e.g. [13–15, 31] and references therein),

M. Maggioni (✉)
Duke University, Durham, NC, USA
e-mail: mauro@math.duke.edu

nonlinear dimension reduction [4, 6, 10, 29, 43, 44, 60, 76, 78, 79, 86–88, 90] and metric space embeddings [10, 12, 59]. Oftentimes, one is interested in studying functions on the data, for the purpose of denoising, fitting, and prediction. These questions can be studied through approximation theory (e.g. [8, 9] and references therein), machine learning [4], and signal processing [34, 70], at least in low-dimensional Euclidean spaces. The combination of the study of geometric properties with the study of functions defined on the data is quite a recent and promising trend [4, 6, 29, 30, 60, 83, 92].

We will for the moment restrict our attention to data sets represented as discrete sets in \mathbb{R}^D . A feature common to many data sets is their high-dimension D , which may range from 10 to 10^6 . This implies that classical statistics, by which we mean the analysis in the case where the dimension D is fixed and the number of points n goes to infinity (or, at least $n \gg 2^D$), is not applicable. Typical situations that we consider have n of the same order as D , and oftentimes $n < D$. In this regime, more appropriate asymptotics are those with n fixed and D going to infinity.

A key observation is that in several situations the data seems to be concentrated along low-dimensional sets in \mathbb{R}^D (e.g. [4, 6, 10, 29, 43, 44, 60, 76, 78, 79, 86–88, 90]). In this case, it is natural to ask what geometric properties these low-dimensional sets have, and how to exploit this phenomenon in order to better model and learn from data.

The interplay between geometry of sets, function spaces on sets, and operators on sets is of course classical in Harmonic Analysis.

This paper gives an overview of very recent work in the geometric analysis of high-dimensional point clouds and tries to briefly summarize the papers [2, 24, 65]. Material related to this paper is available at <http://www.math.duke.edu/~mauro>.

10.2 Multiscale SVD

The quantitative study of geometric properties of sets, such as rectifiability and harmonic analysis, is classical [35–41, 56]. The applications of ideas from geometric measure theory to the analysis of point clouds are emerging, and here we would like to review a small number of very recent ones. One of the basic tools in the analysis of data sets in statistics is Principle Component Analysis (PCA), which is based on the Singular Value Decomposition (SVD). Any $n \times D$ matrix X may be decomposed as $X = U\Sigma V^T$, where $U \in \mathbb{R}^{n \times n}$ and $V \in \mathbb{R}^{D \times D}$ are orthonormal and $\Sigma \in \mathbb{R}^{n \times D}$ is diagonal and positive semidefinite. The diagonal entries $\{\lambda_i\}$ of Σ , called singular values (S.V.'s), are ordered in decreasing order: $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{n \wedge D} \geq 0$. This decomposition is called the SVD of X . It is useful to notice that the first d columns of V span a linear subspace that minimizes $\sum_{i=1}^n \|x_i - P_\Pi(x_i)\|_{\mathbb{R}^D}^2$ over all choices of d -dimensional linear subspaces Π (here P_Π denotes the orthogonal projection onto Π). We say that the first d columns of V produce the d -dimensional least squares fit to X . If the rows $\{x_i\}_{i=1}^n$ of X represent n data points in \mathbb{R}^D , Principle Component Analysis consists of computing the empirical mean $m(X) = \frac{1}{n} \sum_{i=1}^n x_i$, considering

the new matrix \bar{X} whose rows are $\frac{1}{\sqrt{n}}(x_i - m(X))$, and computing the SVD of \bar{X} . The columns of V are called the principal vectors. An alternative interpretation is the following: if we let

$$\begin{aligned} \text{cov}(X) &= \frac{1}{n} \sum_{i=1}^n (x_i - m(X)) \otimes (x_i - m(X)) \\ &= \frac{1}{n} \sum_{i=1}^n (x_i - m(X))^T (x_i - m(X)) \end{aligned} \quad (10.1)$$

be the empirical covariance matrix of X , then from $\bar{X} = U \Sigma V^T$ we immediately deduce that

$$\text{cov}(X) = V \Sigma^T \Sigma V^T \quad (10.2)$$

so that λ_i^2 are the eigenvalues of the covariance matrix. Of course in this setting one thinks of having a random variable X taking values in \mathbb{R}^D , and the mean and covariance are its first and second order statistics. It is clear how to generalize PCA and SVD to the case of infinitely many points.

If the data points $\{x_i\}$ lie, say, uniformly on a bounded domain in a d -dimensional linear subspace, then for n large enough (in fact, $n \gtrsim d \log d$ is enough [77]¹), X will be of rank exactly d . The first d empirical singular values will have the same order as the true ones (i.e. the ones obtained as $n \rightarrow \infty$) and the remaining ones will be exactly 0, and the first d columns of V will span the linear subspace. Because of the least squares fit property, we expect that if we add “small” noise to the points $\{x_i\}$, the smallest singular values should be perturbed by a small amount and would still be much smaller than the top d singular values, indicating the existence of a d -dimensional linear subspace where the (noiseless) data lies.

We are interested in the case where the data points lie on a d -dimensional non-linear manifold \mathcal{M} embedded in a high-dimensional space \mathbb{R}^D and are corrupted by high-dimensional noise. This model has gained popularity in the machine learning community in recent years [4, 6, 10, 29, 43, 44, 60, 76, 78, 79, 86–88, 90]. While only in particular cases may one expect this model to be correct for real data, it is a step beyond linear models. More general models may also be considered, such as unions of manifolds, possibly intersecting each other, and possibly of different dimensions. Understanding these more general models seems to require the understanding of the simpler model with one manifold only. A particular choice of this general model is that of a union of planes, possibly of different dimensions [1, 82].

In order to study these more complicated models, it may seem tempting to discard SVD, which is so well-adapted to linear models. The SVD of a matrix X representing data on a d -dimensional manifold \mathcal{M} in \mathbb{R}^D may not reveal d at all. As a first trivial example, consider a planar circle ($d = 1$) of radius r embedded in \mathbb{R}^D : $\text{cov}(X)$ has exactly 2 nonzero eigenvalues equal to $\frac{r}{\sqrt{2}}$. More generally, it is easy to

¹ These bounds are continuously being refined, in particular in rather general cases the $\log d$ term may be reduced to $\log \log d$.

construct a one-dimensional manifold ($d = 1$) such that $\text{cov}(X)$ has full rank: it is enough to pick a curve that spirals out in more and more dimensions. A simple construction (sometimes called Y. Meyer's staircase) is the following: pick the D points $0, 1, \dots, D-1$ on the real line, and let $\chi_{[0,2)}(x) = 1$ if $x \in [0, 2)$ and 0 otherwise. Then the set

$$\{x_t := \chi_{[0,2)}(\cdot - t)\}_{t \in \mathbb{R}} \subset L^1(\mathbb{R}) \quad (10.3)$$

is a one-dimensional manifold. This may be discretized by evaluating the functions x_t on the discrete set $\{0, 1, \dots, D-1\}$. Notice that x_{t_1} and x_{t_2} are orthogonal whenever $|t_1 - t_2| > 2$, so this curve spirals into larger and larger subspaces as t increases. Similar considerations would hold after discretization of the space and restriction of t to a bounded interval.

However, one may still make good use of PCA if one performs it locally at multiple scales: for every $r > 0$ and every $z \in \mathcal{M}$ consider

$$X_{z,r} := \mathcal{M} \cap B_z(r) \quad (10.4)$$

i.e. the intersection of \mathcal{M} with a Euclidean ball (in the ambient space \mathbb{R}^D) centered at z of radius r . Perform PCA on $X_{z,r}$, and let $\{\lambda_{i,z,r}\}_{i=1}^D$ be the corresponding singular values. Also, let

$$\Delta_i(X_{z,r}) = \lambda_{i,z,r}^2 - \lambda_{i+1,z,r}^2 \quad \text{for } 1 \leq i \leq D-1 \quad (10.5)$$

and $\Delta_D(X_{z,r}) = \lambda_{D,z,r}^2$: these are the gaps of the squared singular values of $X_{z,r}$. For a fixed z , how do these singular values behave? We expect that for small r the top d singular values $\lambda_{1,z,r}, \dots, \lambda_{d,z,r}$ will grow linearly in r and be large compared to the remaining ones, which are associated with normal directions and grow quadratically in r . The principal components corresponding to the top d singular values will approximate the tangent space to \mathcal{M} at z . This allows one to estimate d . As r grows, however, \mathcal{M} will start curving inside $B_z(r)$ and a subset of the bottom singular values will start to grow, eventually (in general) becoming as large as the top d singular values, as in the examples mentioned above. Therefore, the curvature of \mathcal{M} inside \mathbb{R}^D puts an upper bound on the set of scales that may be used to detect d via SVD.

This clear picture becomes more complicated if we add two factors crucial in applications: sampling and noise. We only have a finite number of samples X_n from \mathcal{M} , which will put a lower bound on the values of r : if r is too small, $X_{n,z,r} := X_n \cap B_z(r)$ will simply be empty, or may not contain enough points to be able to determine the intrinsic dimension d . If high-dimensional noise is added to the samples, so that our observations are in the form $x_i + \sigma \eta_i$, with $x_i \in \mathcal{M}$ and $\sigma \eta_i$ representing noise (e.g. $\eta \sim \mathcal{N}(0, I_D)$, with \mathcal{N} denoting the Gaussian distribution), then another lower bound on r arises: if r is small compared to the “size” of the noise $\sigma \eta$, even if we have enough samples in $X_{n,z,r}$, these samples will look high-dimensional because they are scattered in \mathbb{R}^D by the noise. It is only at scales r higher than the “size” of the noise that there is a chance for the top singular values to detect linear growth (in r) because the SVD will detect a noisy tangent plane. But once again, at larger scales

curvature will take over. Of course, in the range of scales above that of the noise and below that dictated by curvature, $X_{n,z,r}$ must have enough samples so that the SVD may be computed reliably. We discuss this problem in some detail in Sect. 10.3, which is a short summary of the work in [65] (see also [63, 64]).

In Sect. 10.4, we discuss the problem of efficiently representing data on a non-linear d -dimensional manifold \mathcal{M} embedded in \mathbb{R}^D , with $d \ll D$, by constructing geometric dictionaries. If \mathcal{M} were linear, then we could perform SVD, use dD numbers to encode the d -dimensional subspace \mathcal{M} lies on (e.g. by the first d columns of the matrix V), and then representing every point on \mathcal{M} would require only d coefficients, instead of D . When \mathcal{M} is not linear nor contained in a low-dimensional subspace of \mathbb{R}^D , it is not clear how to generalize such a construction in order to efficiently store the data. We briefly discuss recent work based on so-called Geometric Wavelets [2, 25], which aims at efficiently encoding the multiscale family of SVD's discussed above by encoding the difference between approximate tangent planes at different scales. This encoding not only reduces the cost of encoding these planes, but yields a multiscale decomposition of every point of \mathcal{M} , and therefore of \mathcal{M} itself, and fast but nonlinear algorithms for computing a geometric wavelet transform and its inverse for every point. This may be thought of as a geometric version of wavelets. Much needs to be explored in these directions. In any case, this yields multiscale matrix decompositions that allow one to efficiently encode the data and reveal structures in the data. Section 10.4 is a short summary of the work in [2] (see also [25]).

Finally, when data is modeled as lying on multiple planes, of possibly different dimensions, and possibly intersecting, we discuss the problem of estimating the family of these planes. This is the topic of Sect. 10.5, which is a short summary of the work [24].

10.3 Intrinsic Dimension Estimation

The problem of estimating the intrinsic dimension of a point cloud is of interest in a wide variety of situations, such as estimating the number of variables in a linear model in statistics, the number of degrees of freedom in a dynamical system, the intrinsic dimension of a data set modeled by a probability distribution highly concentrated around a low-dimensional manifold. Many applications and algorithms crucially rely on the estimation of the number of components in the data, for example spectrometry, signal processing, genomics and economics, to name only a few. Moreover, many manifold learning algorithms [4, 6, 10, 29, 43, 44, 60, 76, 78, 79, 86–88, 90] assume that the intrinsic dimension is given.

When the data lies on a plane – for example, when it is generated by a multivariate linear model – Principle Component Analysis allows one to recover the plane. This case is well understood as the number of samples grows to infinity and also when noise is present (see e.g., out of many works, [54, 74, 81] and references therein).

The finite sample situation is less well understood. Even in the case of points lying on a plane we derive new results using a new approach; however, the situation we are really interested in is that of data having a geometric structure more complicated than linear, such as when the data lies on a low-dimensional manifold. Several algorithms have been proposed to estimate intrinsic dimension in this setting; for lack of space we cite only [11, 14, 16, 18, 32, 47, 49, 50, 53, 58, 61, 71, 75, 85].

10.3.1 Multiscale Dimension Estimation

We start by describing a stochastic geometric model generating the point clouds we will study. Let (\mathcal{M}, g) be a compact, smooth, d -dimensional Riemannian manifold, isometrically embedded in \mathbb{R}^D . Let η be \mathbb{R}^D -valued with $\mathbb{E}[\eta] = 0$, $\text{Var}[\eta] = 1$ (the “noise”), for example, $\eta \sim \mathcal{N}(0, I_D)$. Let $X = \{x_i\}_{i=1}^n$ be a set of uniform (with respect to the natural volume measure on \mathcal{M}), independent random samples on \mathcal{M} . Our observations \tilde{X} are noisy samples: $\tilde{X} = \{x_i + \sigma \eta_i\}_{i=1}^n$, where η_i are i.i.d. samples from η and $\sigma > 0$. These points may also be thought of as being sampled from a probability distribution $\tilde{\mathcal{M}}$ supported in \mathbb{R}^D and concentrated around \mathcal{M} . Here and in what follows we represent a set of n points in \mathbb{R}^D by an $n \times D$ matrix, whose (i, j) entry is the j th coordinate of the i th point. In particular X and \tilde{X} are used to denote both the point cloud and the associated $n \times D$ matrices, and N is the noise matrix of the $\sigma \eta_i$'s.

The problem we concern ourselves with is to **estimate** $d = \dim \mathcal{M}$, **given** \tilde{X} . We shall use multiscale SVD, as described above, and start with an example.

10.3.1.1 Example: d -Dimensional Sphere in \mathbb{R}^D , with Noise

Let $\mathbb{S}^d = \{x \in \mathbb{R}^{d+1} : \|x\|_2 = 1\}$ be the unit sphere in \mathbb{R}^{d+1} , so $\dim(\mathbb{S}^d) = d$. We embed \mathbb{S}^d in \mathbb{R}^D via the natural embedding of \mathbb{R}^{d+1} in \mathbb{R}^D via the first $d+1$ coordinates. We obtain X by sampling n points uniformly at random from \mathbb{S}^d , and \tilde{X} is obtained by adding D -dimensional white Gaussian noise of standard deviation σ in every direction. We call this data set $\mathbb{S}^d(D, n, \sigma)$; Fig. 10.1 plots the multiscale S.V.'s of $\mathbb{S}^9(100, 1000, 0.01)$.

In Fig. 10.2, we consider the multiscale S.V.'s of $\mathbb{S}^9(100, 1000, 0.1)$, as a function of r . Several observations are in order. Firstly, notice that \mathbb{R}^{d+1} is divided into 2^{d+1} sectors, and therefore by sampling 1,000 points on \mathbb{S}^9 we obtain about 1 point per sector (!). Secondly, observe that the noise size, if measured by $\|x_i - \tilde{x}_i\|_2^2$, i.e. by how much each point is displaced, would be of order $\mathbb{E} \sim 1$, which is comparable with the radius of the sphere itself (!). Therefore, this data set may be described as randomly sampling one point per sector at distance 1 from the origin in the first $d+1$ coordinates, then moved by 1 in a random direction in \mathbb{R}^{100} . The situation may seem hopeless.

In fact, we can detect reliably the intrinsic dimension of \mathcal{M} . At very small scales, $B_z(r)$ is empty or contains less than $O(d)$ points, and the rank of $\text{cov}(X_{z,r})$ is even

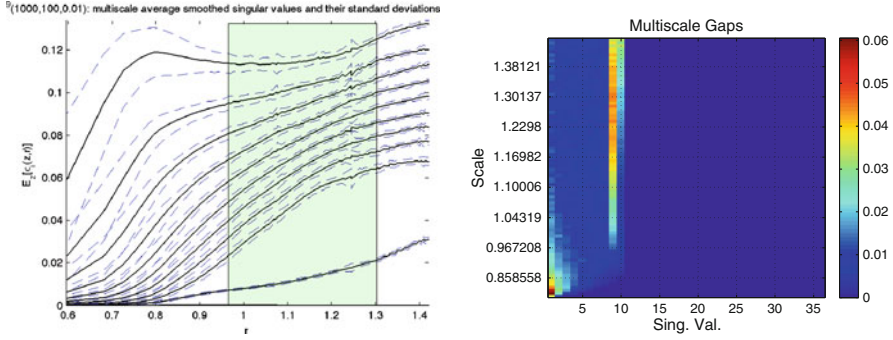


Fig. 10.1 $\mathbb{S}^9(100,1000,0.01)$. *Left*: plot of $\mathbb{E}_z[\lambda_{i,z,r}]$, and corresponding standard deviation bands (dotted), as a function of r . The top 9 S.V.'s dominate and correspond to the intrinsic dimensions; the 10th S.V. corresponds to curvature, and slowly increases with scale (note that at large scales $\Delta_{10} > \Delta_9$); the remaining S.V.'s (not shown) correspond to noise in the remaining 90 dimensions, and converge to the one-dimensional noise size σ . *Right*: plot of the multiscale gaps: on the horizontal axis we have the index i of the gap, and on the vertical axis the scale r . The entry (i, r) is the average (over z) gap $\mathbb{E}_z[\Delta_i(X_{z,r})] := \mathbb{E}_z[\lambda_i(X_{z,r}) - \lambda_{i+1}(X_{z,r})]$. At small scales the noise creates the gaps at the bottom left of the figure; at larger scales we see a large gap at $i = 9$, and at even larger scales that gap is surpassed by the gap corresponding to $i = 10$. This plane is a sort of geometric scale-“frequency” plane, where “frequency” is the index of the singular values




Fig. 10.2 Same as above, but for $\mathbb{S}^9(100,1000,0.1)$, i.e. 10 times larger noise

less than d . At small scales, no gap among the $\lambda_{i,z,r}$ is visible: $B_z(r)$ contains too few points, scattered in all directions by the noise, and new increasing S.V.'s keep arising for several scales. At larger scales, the top $d = 9$ S.V.'s start to separate from the others: at these scales the noisy tangent space is detected. At even larger scales, the curvature starts affecting the covariance, as indicated by the slowly growing 10th S.V., while the remaining smaller S.V.'s tend approximately to the *one-dimensional* noise variance: this is the size of the noise relevant in our procedure, rather than the much larger expected displacement measured in the full \mathbb{R}^D , which was of size $O(1)$.

Motivated by applications to large data sets in high-dimensional spaces that are assumed to be intrinsically low-dimensional, we are interested in the regime where

D is large, $d \ll D$, and we will ask how large n needs to be in order to estimate d correctly with high probability (w.h.p.). In a classical statistical framework, one is often interested in the regime where D, d are fixed and n tends to infinity, but in that case one would conduct the analysis as $r \rightarrow 0$ and this would lead essentially to the problem of consistency of PCA, and noise would be a relatively minor complication. In many applications, D is large and n cannot be taken much larger than D itself: we will therefore be interested in the regime when D and n are large but $\frac{n}{D} = O(1)$.

10.3.2 Results

In the above setting, we are interested in non-asymptotic results that hold for finite n, d , and D , since they will imply finite sample inequalities w.h.p.

We fix a center z for our computations. Let T_z and N_z be the tangent and normal planes to \mathcal{M} at z . Their dimensions are of course d and $D - d$. Let $X_{z,r}^{\parallel}$ and N^{\parallel} be the projections of $X_{z,r}$ and N onto T_z , and let $X_{z,r}^{\perp}$ and N^{\perp} be the projections onto N_z , so that $X_{z,r} = X_{z,r}^{\parallel} + X_{z,r}^{\perp}$ and $N = \sigma_{\parallel} N^{\parallel} + \sigma_{\perp} N^{\perp}$. We assume that for each fixed z , there exist parameters $R_{\max}, R_{\min}, \lambda_{\max}, \lambda_{\min}, \kappa, v_{\min}$, and v_{\max} (which in general will depend on z) such that for every $r \in [R_{\min}, R_{\max}]$:

$$\begin{aligned} \text{spec}(\text{cov}(X_{z,r}^{\parallel})) &\subseteq d^{-1}r^2[\lambda_{\min}^2, \lambda_{\max}^2] \\ \|\text{cov}(X_{z,r}^{\perp})\|_F &\leq \kappa^2 r^4 \\ \text{vol}(X_{z,r}) &\in [v_{\min}, v_{\max}] \mu_{\mathbb{R}^d}(\mathbb{B}^d) r^d, \end{aligned} \quad (10.6)$$

where $\text{spec}(\text{cov}(X_{z,r}^{\parallel}))$ denotes the set of eigenvalues of $\text{cov}(X_{z,r}^{\parallel})$ and \mathbb{B}^d is the Euclidean unit ball in \mathbb{R}^d .

We would like to detect the unknown intrinsic dimension d by estimating a range of “good” scales where the d th gap is the largest gap. We define

$$\tilde{\Lambda}_{z,r} := \{r > 0 : \Delta_d(\tilde{X}_{z,r}) = \max_{i=1,\dots,D} \Delta_i(\tilde{X}_{z,r})\} \quad (10.7)$$

Observe that $\tilde{\Lambda}_{z,r}$ is a random set, and we are interested in finding an interval which is contained in $\tilde{\Lambda}_{z,r}$ with high probability given n noisy samples as above.

Proposition 10.1 ($n \rightarrow \infty$). Assume $\lambda_{\max} = \lambda_{\min} = \lambda$, $r > \sigma_{\parallel} \sqrt{d} + \sigma_{\perp} \sqrt{D}$. Then for n large enough, a sufficient condition for $r \in [R_{\min}, R_{\max}]$ being in $\tilde{\Lambda}_{z,r}$ is that:

$$\underbrace{\frac{\lambda^2 r^2}{d}}_{\text{tangent term}} + \underbrace{\sigma_{\parallel}^2}_{\text{tangent noise}} \geq \underbrace{2\kappa^2 r^4 + \frac{\lambda \kappa r^3}{\sqrt{d}} \vee \sigma_{\perp}^2}_{\text{curvature term}} + \underbrace{\sigma_{\perp}^2}_{\text{normal noise}} + O(n^{-\frac{1}{2}}) \quad (10.8)$$

As $\sigma_{\parallel}, \sigma_{\perp} \rightarrow 0$, this is implied by $r \leq \frac{\lambda}{2\kappa\sqrt{d}}$.

Proposition 10.2 ($n, D \rightarrow \infty, \frac{n}{D} \rightarrow \gamma, \sigma_{\perp} = \frac{\sigma}{\sqrt{D}}$). Assume $\lambda_{\max} = \lambda_{\min} = \lambda$. Then for n, D large enough, and $\gamma = \frac{n}{D}$, a sufficient condition for $r \in [R_{\min}, R_{\max}]$ being in $\tilde{\Lambda}_{z,r}$ is that:

$$\underbrace{(\sigma_{\parallel} \sqrt{d} + \sigma)}_{\text{noise}} \lesssim r \lesssim \underbrace{\frac{\lambda}{\kappa \sqrt{d}}}_{\text{curvature}} \quad (10.9)$$

Proposition 10.3 ($D \rightarrow \infty, \sigma_{\perp} = \frac{\sigma}{\sqrt{D}}$). Assume $\lambda = \lambda_{\max} = \lambda_{\min}$ and c is a universal constant. A sufficient condition for $r \in [R_{\min}, R_{\max}]$ being in $\tilde{\Lambda}_{z,r}$ is that:

$$\underbrace{(\sigma_{\parallel} \sqrt{d} + \sigma)}_{\text{noise}} \vee \underbrace{\left(\frac{d \log(d)}{n} \frac{c \operatorname{vol}(\mathcal{M})}{\lambda^2 v_{\min} \mu_{\mathbb{R}^d}(\mathbb{B}^d)} \right)^{\frac{1}{d}}}_{\text{sampling}} \lesssim r \lesssim \underbrace{\frac{\lambda}{\kappa \sqrt{d}}}_{\text{curvature}} \quad (10.10)$$

In fact, in all of the above results, the conditions are sufficient not only in the limit, but with high probability for finite values of n, D . These results are more technical and the interested reader is referred to [65]. They essentially imply, in this context, that as soon as $n_r := |X_{z,r}| \gtrsim_{\kappa, \sigma_{\parallel}, \sigma_{\perp}, v_{\min}} d \log d$ for $r < R_{\max}$, then $\tilde{\Lambda}_{z,r}$ is non-empty with high-probability. An efficient algorithm for finding r 's in $\tilde{\Lambda}_{z,r}$ is also developed in [65], and tested against the leading competitors (see the following section). Finally, the setting in [65] is much more general than the one presented above; in particular no manifold assumption is made. Instead, the existence of a set of scales is assumed, at which the data set looks d -dimensional plus smaller detail structure and noise.

In the special case of \mathbb{S}^d , if we have no noise and $n \rightarrow \infty$ one can show that:

$$\begin{aligned} \lambda_{i,z,r}^2 &= \frac{1}{d+2} r^2 + O(r^4) & \text{for } 1 \leq i \leq d \\ \lambda_{d+1,z,r}^2 &= \frac{d}{(d+2)^2(d+4)} r^4 + O(r^6) \\ \lambda_{i,z,r}^2 &= 0 & \text{for } i > d+1 \end{aligned}$$

So here, $\lambda_{\max}^2 = \lambda_{\min}^2 = \frac{d}{d+2} \sim 1$ and $\kappa \sim \frac{1}{d}$. Although on first glance it appears that Proposition 10.1 gives us that Δ_d is the largest gap when $r \lesssim \sqrt{d}$, this is in fact not the case since R_{\max} (the upper bound on the region where the curvature S.V. $\lambda_{d+1,z,r}$ grows quadratically with respect to the tangent S.V.'s) is actually small: $R_{\max} = O(1)$.

In all of these results, a “curse of intrinsic dimension” is visible. By “curse of dimensionality” one usually means the large number of samples needed for estimating functions of many parameters. For example, if one tries to approximate a continuous function on the d -dimensional unit cube \mathbb{Q}^d up to precision ε , one needs in general one sample in every little d -dimensional cube of side ε contained in \mathbb{Q}^d :

the number of such cubes is ε^{-d} , which is large as soon as ε is not large and d is not small (for example: if $\varepsilon = 10^{-3}$ and $d = d$, one would need 10^{3d} samples to approximate the function up to only 3 digits). From the geometric perspective, the curse of dimensionality may manifest itself in terms of concentration of measure phenomena. In our particular situation, for example, covariance matrices of intrinsically high-dimensional objects tend to be small, and therefore easily corrupted by noise. For example, the covariance matrix of the $(d-1)$ -dimensional unit sphere \mathbb{S}^{d-1} is $\frac{1}{d}I_d$ (and not I_d as one may have expected). In particular, if Gaussian noise $\sigma\mathcal{N}(0, I_d)$ is added to points sampled on \mathbb{S}^{d-1} , then the covariance “signal to noise ratio” is $d^{-\frac{1}{2}}/\sigma$, which goes to 0 as $d \rightarrow +\infty$. In the last two Propositions, one sees this curse in the upper bound on the right hand side, which contains the factor $d^{-\frac{1}{2}}$. However, notice that if $\kappa \sim d^{-\frac{1}{2}}$, such bounds become independent of d . In other words, the curse of intrinsic dimensionality, in this context, is not only rather mild, but disappears by decreasing the curvature κ as the intrinsic dimension increases. This is an analogue of sort to assuming smoothness dependent on d to break the curse in approximating functions in high-dimensions, an interesting approach taken in functional analysis, approximation theory and statistics.

10.3.3 Algorithm

The results above suggest the following algorithm: for each $z \in \mathcal{M}$, $r > 0$, $i = 1, \dots, D$, we compute $\lambda_{i,z,r}$. When r is large, if \mathcal{M} is contained in a linear subspace of dimension K we will observe K large eigenvalues and $D-K$ smaller noise eigenvalues, in the regime for the values of K, D, σ, n suggested by our results. Clearly, $d \leq K$. Moreover, $\{\lambda_{i,z,r}\}_{i=K+1, \dots, D}$ will be highly concentrated and we use them to estimate σ , which is useful per se. By viewing $\{\lambda_{i,z,r}\}_{i=K+1, \dots, D}$, we identify an interval in r where the noise is almost flat, i.e. we remove the small scales where the distortion due to noise dominates.

We look at the first $\{\lambda_{i,z,r}\}_{i=1, \dots, K}$, and the goal is to decide how many of them are due to the extrinsic curvature of \mathcal{M} . But the curvature S.V.’s grow quadratically w.r.t. the “tangential” (non-curvature) S.V.’s: a best least-square linear and quadratic fit to $\lambda_{i,z,r}$, as a function of r , is enough to tell the curvature S.V.’s from the tangential S.V.’s. See Fig. 10.3 for pseudo-code of the described algorithm.

MATLAB code and a User Interface for navigating the multiscale S.V.’s are available at www.math.duke.edu/~mauro.

10.3.4 Examples

We test our algorithm on several data sets obtained by sampling manifolds, and compare it with existing algorithms. The test is conducted as follows. We fix the


```

 $[\hat{d}, \hat{R}_{\min}, \hat{R}_{\max}] = \text{EstDimMSVD}(\tilde{X}_n, K)$ 

// Input:
//  $\tilde{X}_n$  : an  $n \times D$  set of noisy samples
//  $K$  : upper bound on the intrinsic dimension  $d$ 

// Output:
//  $\hat{d}$  : estimated intrinsic dimension
//  $(\hat{R}_{\min}, \hat{R}_{\max})$  : estimated interval of good scales

Nets = MultiscaleNets( $\tilde{X}_n, K$ )
 $\lambda_{K+1,z,r} = \text{FindLargestNoiseSingularValue}(\tilde{X}_n, \text{Nets})$ 
 $\hat{R}_{\min} = \text{Smallest scale for which } \lambda_{K+1,z,r} \text{ is decreasing and } |B_z(\hat{R}_{\min})| \gtrsim K \log K$ 
 $\hat{R}_{\max} = \text{Largest scale for which } \lambda_{1,z,r} \text{ is nonincreasing}$ 
 $\hat{d} = \text{Largest } i \text{ such that:}$ 

- for  $r \in (\hat{R}_{\min}, \hat{R}_{\max})$ ,  $\lambda_{i,z,r}$  is linear and  $\lambda_{i+1,z,r}$  is quadratic in  $r$ , and
- $\Delta_i^{(z,r)}$  is largest gap for  $r$  in a large fraction of  $(\hat{R}_{\min}, \hat{R}_{\max})$

 $(\hat{R}_{\min}, \hat{R}_{\max}) = \text{Largest interval in which } \Delta_{\hat{d}}^{(z,r)}$  is the largest gap

```

Fig. 10.3 Pseudo-code for the intrinsic dimension estimator based on multiscale SVD

ambient space dimension to $D = 100$. We let \mathbb{Q}^d , \mathbb{S}^d , \mathcal{S} , \mathcal{Z}^d be, respectively, the unit d -dimensional cube, the unit d -dimensional sphere, a manifold product of an S -shaped curve of roughly unit diameter and a unit interval, and the d -dimensional Meyer’s staircase. Each of these manifolds is embedded isometrically in \mathbb{R}^K , where $K = d$ for \mathbb{Q}^d , $K = d + 1$ for \mathbb{S}^d , $K = 3$ for \mathcal{S} and $K = D$ for \mathcal{Z}^d , and \mathbb{R}^K is embedded naturally in \mathbb{R}^D . Finally, a random rotation is applied (this is irrelevant since all the algorithms considered are invariant under isometries). We draw n samples uniformly (with respect to the volume measure) at random from each manifold, and add noise $\sigma \eta \sim \sigma \mathcal{N}(0, I_D)$. We incorporate these parameters in the notation by using $\mathbb{Q}^d(n, \sigma)$ to denote the set of n samples obtained as described above, where the manifold is the d dimensional unit cube and the noise has variance σ^2 ; if $D \neq 100$, we write $\mathbb{Q}^d(D, n, \sigma)$, similarly for \mathbb{S}^d , \mathcal{S} , \mathcal{Z}^d . We also consider a variation of these sets, where we dilate \mathbb{R}^K (after embedding the manifold but before any other operation) by a diagonal dilation with factors drawn at random from the multiset $\{1, 1, 1, 1, 0.9, 0.9, 0.9, 0.8, 0.8\}$. See Fig. 10.4 and Fig. 10.5 for plots of the multiscale S.V.’s of $\mathcal{S}(1000, 0.01)$ and $\mathcal{Z}^{20}(500, 1000, 0.05/\sqrt{1000})$.

We consider here $d = 6, 12, 24, 48$ for \mathbb{Q}^d and \mathbb{S}^{d-1} , $d = 10, 20, 50$ for \mathcal{Z}^d . The samples size is set as $n = 250, 500, 1000, 2000$. We let the noise parameter $\sigma = 0, 0.01, 0.025, 0.05, 0.1, 0.15, 0.2$. For each combination of these parameters we generate 5 realizations of the data set and report the most frequent (integral) dimension returned by the set of algorithms specified below, as well as the standard deviation of the estimated dimension. We test the following algorithms: “Debiasing” [16], “Smoothing” [18], RPM in [51], “MLE” [62], “DeBias” [17], “kNN” [33], “SmoothKNN” [19], as well as the classical Correlation Dimension and Taken

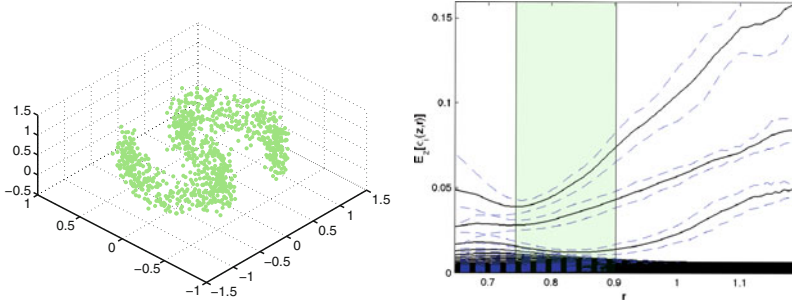


Fig. 10.4 *Left:* The S-manifold $\mathcal{S}(1000, 0.01)$ corrupted by noise. *Right:* its average multiscale singular values. The green bands are the set of good scales returned by the algorithm




Fig. 10.5 The average multiscale singular values of the Meyer staircase $\mathcal{Z}^{20}(500, 1000, 0.05/\sqrt{1,000})$ corrupted by noise. The k th point in this Meyer staircase is a 1,000-dimensional vector whose entries from $k+1$ to $k+20$ are equal to 1, and all other entries are 0. The green bands are the set of good scales returned by the algorithm

estimator [46, 85]. See Figs. 10.6–10.8 and Table 10.1 for a sample of the results. The “MFA” label in Table 10.1 refers to one of the state-of-art Bayesian approaches to dimension estimation [26]. For MFA, the authors of [26] ran the code, given only the information that no data set would have intrinsic dimension larger than 100 (this is essentially the only parameter in our algorithm, and it is used only for speeding up computations); in all the other cases we ran the code ourselves, after finding a reasonable range of the parameters that worked on toy examples. There is a large disparity in the number of parameters in these algorithms, ranging from 1 to 7. We will make the data sets publicly available at www.math.duke.edu/~mauro so that other researchers may try their algorithms (present and future) on a common set of examples.

Finally, in Fig. 10.9 we consider a data set with different dimensionality in different regions, and run the algorithm pointwise. We show both the pointwise estimated dimensionality, and the maximal value of r the algorithm returns as a good scale.

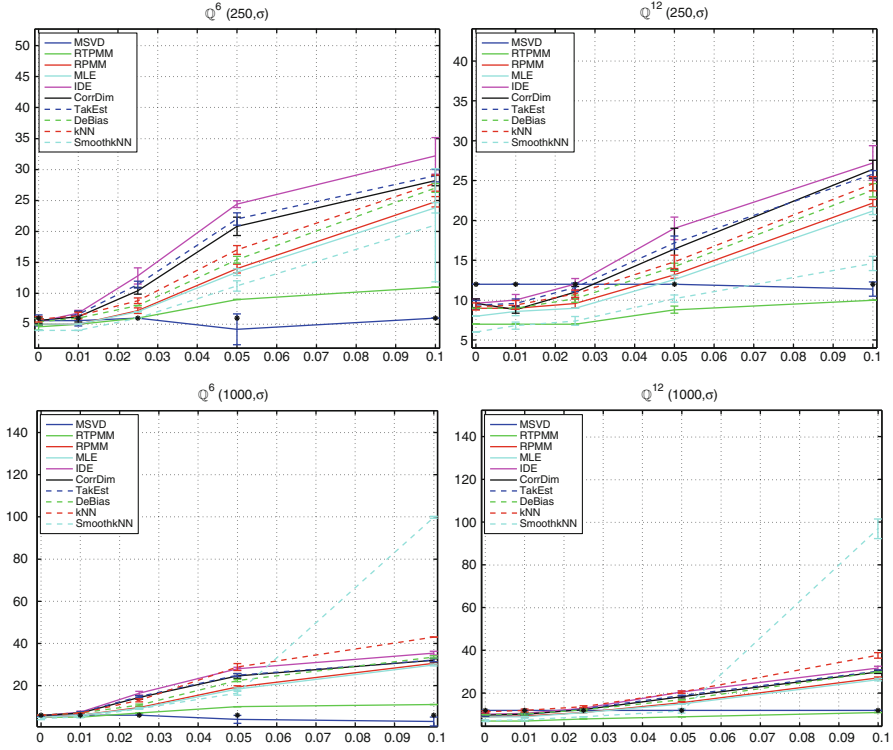


Fig. 10.6 Benchmark data sets: *cube*. The horizontal axis is σ , the one-dimensional standard deviation of the noise, the vertical axis is the estimated dimension. Black dots mark the correct intrinsic dimension

10.4 Geometric Wavelets

Recent work, both in machine learning, signal processing, image analysis, and harmonic analysis, has focused on either the construction of carefully hand-crafted dictionaries for large classes of data sets (wavelet for 1-D signals, curvelets for certain classes of images, chirplets, etc.), or on dictionaries that are tuned to specific data sets [1, 82, 91] (and references therein). The latter approach typically is formulated by requesting to find a dictionary Φ with I elements, such that every element in the data set may be represented, up to a certain precision ε , by at most m elements of the dictionary. This sparsity requirement of the representation is very natural in statistics, signal processing, and interpretation of the representation. Of course, the smaller I and m are, for a given ε , the better. Current constructions of such dictionaries, such as K-SVD [1], k-flats [82], and Bayesian methods [91], cast these requirements as an optimization problem and rely on black-box optimization to find solutions. Typically, no guarantees are provided about the size of I and m (as functions of ε), the computational costs, and the uniqueness of the solution

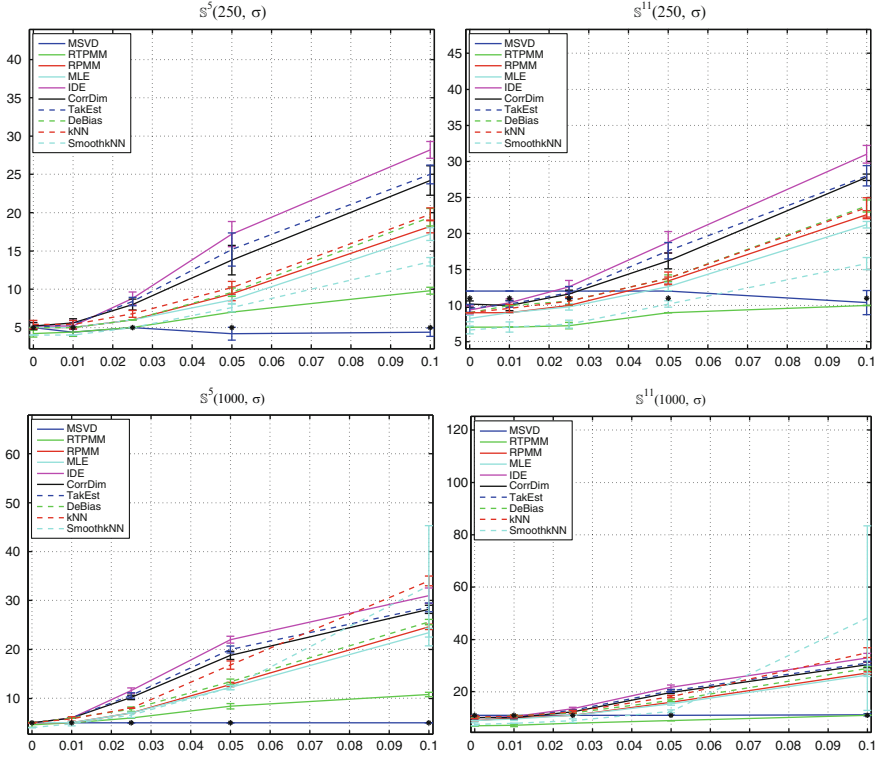


Fig. 10.7 Benchmark data sets: *sphere*. The horizontal axis is σ , the one-dimensional standard deviation of the noise, the vertical axis is the estimated dimension. Black dots mark the correct intrinsic dimension

(in practice, it is observed that these algorithms depend heavily on the initial guess). On the other hand, when these algorithms do give solutions that a practitioner considers “good”, the use of these data-dependent dictionaries can yield very impressive results in such diverse problems as data modeling, classification, image compression and inpainting, and more (e.g. [69] and references therein). Another drawback of existing constructions of data-dependent dictionaries is that the output dictionary is in general highly overcomplete and completely unstructured (even if, at least in the case of images, one may empirically observe certain structures, symmetries and regularities in the dictionary elements). As a consequence, in general there is no fast algorithm for computing the coefficients of the representation of a data point in the dictionary (nor, but less importantly, to sum a long linear combination of dictionary elements), which requires appropriate sparsity-seeking algorithms.

In [24], the authors construct data-dependent dictionaries using a multiscale geometric analysis of the data, based on the geometric analysis in the work [55]. These dictionaries are structured in a multiscale fashion and can be computed efficiently;

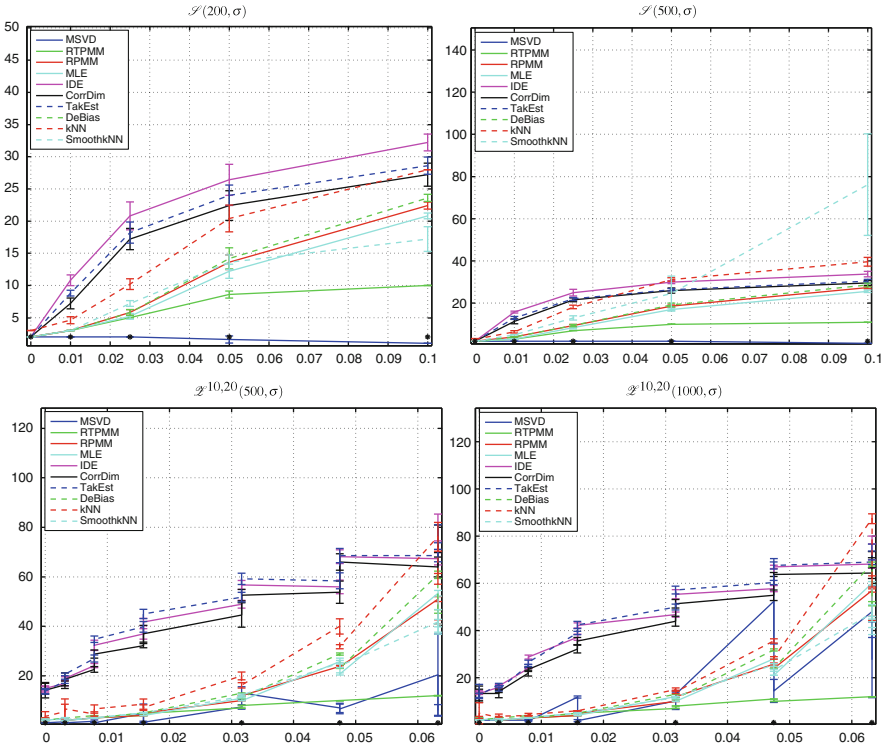


Fig. 10.8 Benchmark data sets: S-shaped manifold \mathcal{S} and Meyer's staircase \mathcal{Z} . The results for \mathcal{Z}^{20} are consistently better than those for \mathcal{Z}^{10} , once the number of points and the level of noise is fixed. This is consistent with the fact that \mathcal{Z}^{20} has a smaller effective curvature than \mathcal{Z}^{10}

the expansion of a data point on the dictionary elements is guaranteed to have a certain degree of sparsity m and can be computed by a fast algorithm; the growth of the number of dictionary elements I (as a function of ϵ) is controlled theoretically and easy to estimate in practice. The elements of these dictionaries are called *Geometric Wavelets* [24], since in some respects they generalize wavelets from vectors that analyze functions to affine vectors that analyze point clouds. The multiscale analysis associated with geometric wavelets shares some similarities with that of standard wavelets (e.g. fast transforms, a version of two-scale relations, etc.), but is in fact quite different in many crucial respects. It is highly nonlinear, as it adapts to arbitrary nonlinear manifolds, albeit every scale-to-scale step is linear (which is key to efficient computation and fast algorithms); translations or dilations do not play any role here, while they are often considered crucial in classical wavelet constructions. Geometric wavelets may allow the design of new algorithms for manipulating point clouds similar to those used for wavelets to manipulate functions. Dictionaries of basis functions have a large number of applications in mathematics and engineering [7, 20, 27, 28, 34, 42, 70, 89].

Table 10.1 This table contains the dimension estimates for a quite benign regime with 1,000 samples and no noise. Even in this straightforward setting the estimation of dimension is challenging for most methods

	RTPMM	RPMM	MLE	IDE	CorrDim	TakEst	DeBias	kNN	SmoothkNN	MFA	MSVD
Q^6	5	5	5	6	5	6	6	6	4	2	6
Q^{12}	7	9	9	10	10	10	10	12	7	4	12
Q^{24}	9	16	16	17	17	17	17	20	11	1	24
Q^{48}	11	26	25	29	28	28	27	32	19	2	48
Q^6	5	5	5	6	6	6	6	6	5	1	6
Q^{12}	7	9	9	10	10	10	10	12	7	2	12
Q^{24}	9	16	15	17	17	17	17	20	11	3	24
Q^{48}	11	26	25	28	27	28	27	31	17	2	48
S^5	4	5	5	5	5	5	5	5	4	2	5
S^{11}	7	9	9	10	10	10	10	10	8	1	11
S^{23}	10	16	16	18	18	18	18	18	14	2	24
S^{47}	11	27	26	31	30	31	29	29	21	3	48
S^5	5	5	5	5	5	5	5	5	4	2	5
S^{11}	7	9	9	10	10	10	10	10	8	1	11
S^{23}	9	16	16	18	18	18	18	18	13	1	23
S^{47}	11	27	26	31	30	30	29	29	21	3	48
\mathcal{S}	2	2	2	2	2	2	2	2	2	1	2
\mathcal{S}	2	2	2	2	2	2	2	3	2	2	2
\mathcal{Z}^1	NaN	NaN	3	340	0	29	3	87	7	4	2
\mathcal{Z}^1	NaN	NaN	2	93	0	14	2	67	3	2	1
\mathcal{Z}^1	NaN	NaN	3	14	12	14	3	3	2	2	1
\mathcal{Z}^1	NaN	NaN	2	13	13	13	2	5	2	2	1

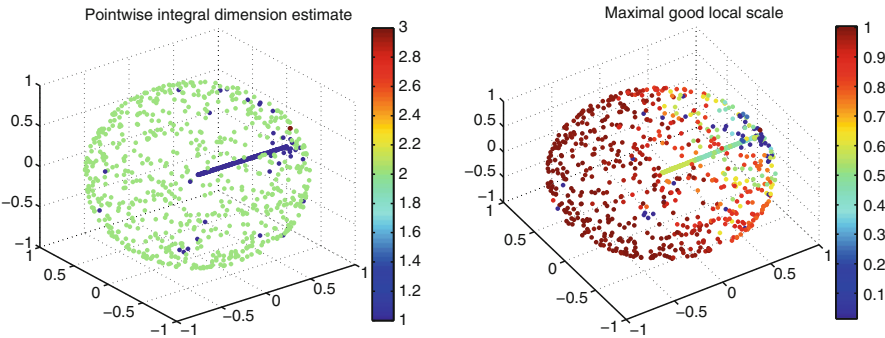


Fig. 10.9 This data set is a union of two manifolds of different dimensions: a 2-dimensional sphere and a 1 dimensional segment. *Left*: the estimated pointwise dimensionality. *Right*: the largest “good” scale returned by the algorithm, for every point. Note how the largest good scale is large for points far from the intersection of the sphere and the segment and decreases as we approach the intersection

10.4.1 Construction of Geometric Wavelets

Let (\mathcal{M}, g) be a d -dimensional compact Riemannian manifold isometrically embedded in \mathbb{R}^D . We are interested in the regime $d \ll D$. Assume we have n samples drawn i.i.d. from \mathcal{M} , according to the natural volume measure $d\text{vol}$ on \mathcal{M} ². We construct a multiscale decomposition of the manifold \mathcal{M} as follows.

We start by decomposing \mathcal{M} into multiscale nested partitions \mathcal{P}_j . For $j \leq J$, let $\mathcal{P}_j = \{C_{j,k}\}_{k \in \Gamma_j}$ be a disjoint cover of \mathcal{M} ; each $C_{j,k}$ contains a ball of radius $\sim 2^{-j}$, has diameter $\sim 2^{-j}$ and piecewise smooth boundary. Moreover, we assume that every $C_{j,k} = \cup_{k' \in \text{children}(j,k)} C_{j+1,k'}$; this also defines $\text{children}(j,k)$. There is a natural tree structure \mathcal{T} associated with this family of partitions. For $x \in \mathcal{M}$, we let $C_{j,x}$ be the unique element of \mathcal{P}_j that contains x . We also note in advance that we will adopt similar notation $(P_{j,x}, \Phi_{j,x}, \Psi_{j,x}, \text{etc.})$ for objects associated with $C_{j,x}$. In practice, we use METIS [57] on a nearest-neighbor weighted graph in order to compute the multiscale partitions \mathcal{P}_j .

For every $C_{j,k}$ we may compute the top d eigenvalues and eigenvectors of the covariance matrix $\text{cov}_{j,k}$ of the data in $C_{j,k}$, as we did before for the matrix of the data in a ball of radius r centered around a point z . Let $\Phi_{j,k}$ be the $D \times d$ orthogonal matrix of the top d eigenvectors of $C_{j,k}$ and $\lambda_1 \geq \dots \geq \lambda_d \geq 0$ the corresponding eigenvalues. We think of the span of the columns of $\Phi_{j,k}$, denoted $\langle \Phi_{j,k} \rangle$, as an approximate tangent space to \mathcal{M} at scale j and location marked by $C_{j,k}$. Let $P_{j,k}$ be the associated affine projection

$$P_{j,k}(x) = \Phi_{j,k} \Phi_{j,k}^* (x - \bar{c}_{j,k}) + \bar{c}_{j,k}. \quad (10.11)$$

where $\bar{c}_{j,k}$ is $\text{avg}(C_{j,k})$. We let, for $\forall x \in \mathcal{M}$ and $j \leq J$,

$$x_{j,J} \equiv P_{\mathcal{M}_{j,J}}(x) := P_{j,x} \circ P_{j+1,x} \circ \dots \circ P_{J,x}(x) \quad (10.12)$$

and define the approximation $\mathcal{M}_{j,J}$ to \mathcal{M} at scale j as the set

$$\mathcal{M}_{j,J} := \cup_{k \in \Gamma_j} P_{\mathcal{M}_{j,J}}(C_{j,k}). \quad (10.13)$$

Although $\mathcal{M}_{j,J}$ and $x_{j,J}$ depend on the finest scale J , we will from now on drop the J subscript for simplicity of notation. \mathcal{M}_j is a coarse approximation of \mathcal{M} at scale j , analogous to what the projection of a function onto a scaling function subspace is in wavelet theory. Under suitable assumptions, $\mathcal{M}_j \rightarrow \mathcal{M}$ in the Hausdorff distance, as $J \rightarrow +\infty$.

The set of pieces of affine planes $P_{j,k}(C_{j,k})$ centered at $\bar{c}_{j,k}$ and spanned by $\Phi_{j,k}$ for all $k \in \Gamma_j$ is an approximation of the manifold \mathcal{M} at scale j , similar to how a scaling function approximation to a function is a coarse version of the function.

² More general hypotheses on the sampling procedure or on the measure μ are possible but we do not consider them here.

We can now construct wavelet planes that span the space needed to complete $\Phi_{j,k}$ into the span of $\{\Phi_{j+1,k'}\}_{k' \in \text{children}(j,k)}$. We leave the description of the construction to [2]. In its simplest form, this construction yields, for every $(j+1, k')$, a $D \times d'_{j+1,k'}$ orthogonal matrix $\Psi_{j+1,k'}$ spanning the subspace $(I - \Phi_{j,k}\Phi_{j,k}^*)\langle \Phi_{j+1,k'} \rangle$. Let $\mathcal{Q}_{j+1,k'}$ be the corresponding affine projection:

$$\mathcal{Q}_{j+1,k'}(x_{j+1}) = \Psi_{j+1,k'}\Psi_{j+1,k'}^*(x_{j+1} - \bar{c}_{j+1,k'}) + (I - \Phi_{j,k}\Phi_{j,k}^*)(\bar{c}_{j,k} - \bar{c}_{j+1,k'}). \quad (10.14)$$

We have the fundamental two-scale relation

$$P_{\mathcal{M}_{j+1}}(x) = P_{\mathcal{M}_j}(x) + \mathcal{Q}_{j+1,x}(P_{\mathcal{M}_{j+1}}(x)) \quad (10.15)$$

for every $x \in \mathbb{C}_{j+1,k'}$. By iterating, we obtain a wavelet sum representation of any point $x \in \mathcal{M}$.

The geometric scaling and wavelet coefficients $\{p_{j,x}\}, \{q_{j,x}\}$ of a point $x \in \mathcal{M}$ are defined by the equations

$$x_j = \Phi_{j,x}p_{j,x} + \bar{c}_{j,x} \quad (10.16)$$

$$\mathcal{Q}_{j+1,x}(x_{j+1}) = \Psi_{j+1,x}q_{j+1,x} + w_{j+1,x} \quad (10.17)$$

$$w_{j+1,x} = (I - \Phi_{j,x}\Phi_{j,x}^*)(\bar{c}_{j,x} - \bar{c}_{j+1,x}), \quad (10.18)$$

where $x_j = P_{j,x}(x)$. The computation of the coefficients (and translations), from fine to coarse, is simple and fast. For any $x \in \mathcal{M}_J$ and $j_0 < J$, the set of coefficients

$$\hat{x} = (q_{J,x}, q_{J-1,x}, \dots, q_{j_0+1,x}, p_{j_0,x}) \in \mathbb{R}^{d + \sum_{j=j_0+1}^J d_{j,x}} \quad (10.19)$$

is called the discrete geometric wavelet transform of x . Since $d_{j,x} \leq d$, $d + \sum_{j>j_0} d_{j,x} \leq (J - j_0 + 1)d$.

Observe that we may immediately extend this transform to points not on \mathcal{M} , but within the so called $\text{reach}(\mathcal{M})$, which is the set of points in the ambient space which have a unique closest point in \mathcal{M} . This set of points may be thought of as a maximal tube, of variable radius, around \mathcal{M} , which does not “self-intersect”.

Finally, one may show that the algorithm for constructing the $\Phi_{j,k}$'s and $\Psi_{j,k}$'s only costs $\mathcal{O}(Dn(\log(n) + d^2 + k))$ [2].

10.4.2 Examples

We conduct numerical experiments in this section to demonstrate the performance of the algorithm.

10.4.2.1 Low-Dimensional Smooth Manifolds

We first consider a simple data set of a 2-dimensional S -shaped manifold in \mathbb{R}^{100} and apply the algorithm to obtain the geometric wavelet transform of the sampled data (3,000 points without noise) in Figs. 10.10 and 10.11. The resulting wavelet coefficients matrix is very sparse (with about 63% of the coefficients below 1% of the maximal magnitude). The reconstructed manifolds also approximate the original manifold well.

10.4.2.2 A Data Set

We next consider a data set of images from the MNIST data set³. We consider the handwritten digit 7. Each image has size 28×28 . We randomly sample 5,000 such images from the database and then project the samples into the first 120 dimensions by SVD. We apply the algorithm to construct the Geometric Wavelets and show the wavelet coefficients at all scales in Fig. 10.12. We observe that the magnitudes of the coefficients stops decaying after a certain scale. This indicates that the data is not on a smooth manifold. We expect optimization of the tree and of the dimension of the wavelet in future work to lead to an efficient representation also in this case.

We then fix two data points (i.e. two images) and show in Figs. 10.13 and 10.14 their reconstructed approximations at all scales and the corresponding wavelet bases (all of which are also images). We see that at every scale we have a handwritten




Fig. 10.10 Geometric wavelet transform of an S -shaped manifold, from which 3,000 points are randomly sampled. *Left*: the reconstructed manifold \mathcal{M}_4 at scale 4. *Right*: the reconstructed manifold \mathcal{M}_8 at scale 8

³ Available, together with detailed description and state-of-art results, at <http://yann.lecun.com/exdb/mnist/>.

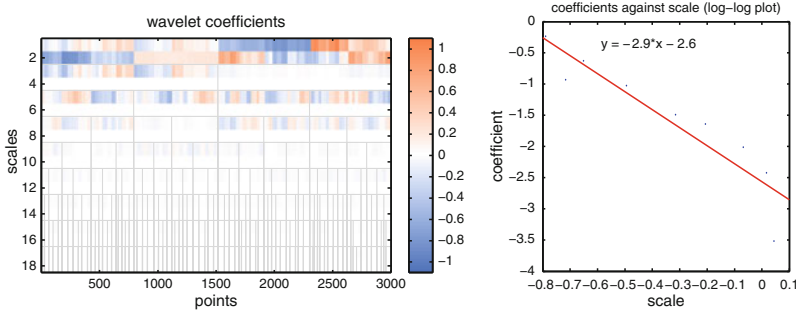


Fig. 10.11 *Left*: Geometric wavelet representation of the data. The x -axis indexes the points (arranged according to the multiscale partition tree), and the y axis indexes the wavelet coefficients as in (10.19), with the coarsest scale at the top and the finest scale at the bottom. The wavelet subspaces have dimension at most 2, and in fact numerically their dimension is, up to two digits of accuracy, 1. This “matrix” is sparse, with about 37% entries above 10^{-2} . Observe that the set of rows is not in one-to-one correspondence with the dictionary elements, since each cell in the tree has its own local dictionary. *Right*: Average error in approximating a point on the manifold, as a function of scale (smaller scales on the right)




Fig. 10.12 Geometric wavelet representation of the data for the images of handwritten 7's. This matrix is less sparse than what we would expect for a manifold. This is an artifact of the construction of geometric wavelets we presented here, in which the dimension of the planes $\langle \Phi_{j,k} \rangle$ is chosen independently of j, k . This constraint is not necessary and is removed in [2], which allows one to tune this dimension, as well as the dimension of the wavelet spaces, to the local (in space and scale) properties of the data

digit, an approximation to the fixed image, and those digits are refined successively to approximate the original data point. The elements of the dictionary quickly fix the orientation and the thickness, and then they add other distinguishing features of the image being approximated.

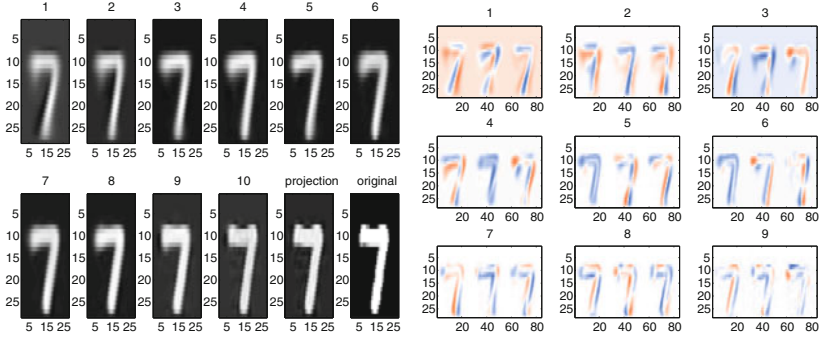


Fig. 10.13 An image representation of the reconstructed data point and the corresponding subset of the wavelet dictionary. *Left*: in images 1–10 we plot coarse-to-fine geometric wavelet approximations of the original data point represented in the last two images (projection and original) on the bottom. *Right*: elements of the wavelet dictionary (ordered from coarse to fine in 1–9) used in the expansion above

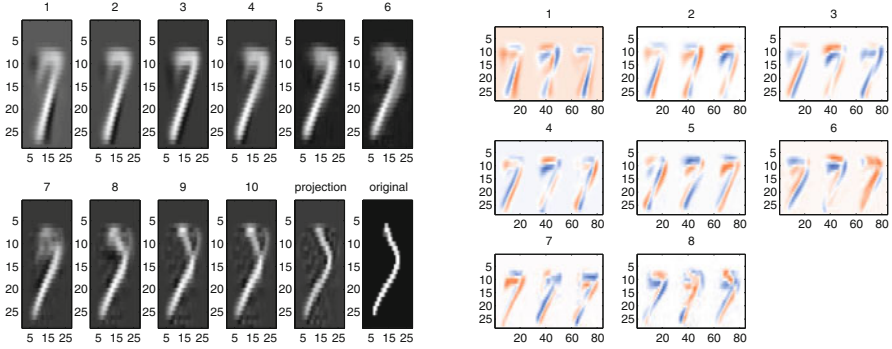


Fig. 10.14 An image representation of the reconstructed data point and the corresponding subset of the wavelet dictionary. *Left*: in images 1–10 we plot coarse-to-fine geometric wavelet approximations of the original data point represented in the last two images (projection and original) on the bottom. *Right*: elements of the wavelet dictionary (ordered from coarse to fine in 1–8) used in the expansion above

10.5 Multiple Planes

Recent work in machine learning, imaging and harmonic analysis has focused on modeling data, or spaces of signals, as a union of linear subspaces (in some cases, affine). More precisely, assume that the data is generated by sampling (uniformly, say) from

$$\bigcup_{i=1}^K \underbrace{a_i(\mathbb{Q}^{d_i})}_{=: \pi_i} \quad (10.20)$$

where the a_i are affine maps $\mathbb{R}^{d_i} \rightarrow \mathbb{R}^D$ that are invertible on their range. In other words, the data is sampled from affine images of unit cubes – pieces of planes – of

different dimensions. These pieces may intersect each other. A typical problem is the following: given n samples, K and d_i , find a_i . Current approaches are based on black-box optimization [1,82], “higher-order” spectral clustering methods [22,23,48] (and references therein), or Monte-Carlo sampling within a Bayesian framework [91]. These cited papers also contain many references to related work. In the first two approaches, K and d_i (or, in some cases, only an upper bound on d_i) are needed as input. In the latter work, experimentation shows that hyper-parameters need to be set very close to the correct values unless the number of points n is very large.

Based on the work on the estimation of intrinsic dimension and determination of “good scales” described in Sect. 10.2 (see [65]), we can tackle this problem not only without having to know K and d_i , but also with guarantees that given a very small number of samples we are going to determine the correct a_i ’s. None of the algorithms described above has this type of guarantee with so little information.

We shall assume that each of the K pieces of planes contains at least c/K points, with high probability, where c is some numerical constant. We shall also assume that a fraction of each piece of plane is far from the intersection with other pieces of plane. The algorithm we propose is the following: First, we pick a sample point x_0 at random. We may assume without loss of generality that x_0 belongs to π_0 . We then run the intrinsic dimension estimator of [65] (see Sect. 10.2). If the estimator succeeds, it will return a range of “good scales”, as well as an estimate of d_0 , which is correct w.h.p.. If the estimator fails, the point was too close to an intersection between π_0 and some other π_i . We simply consider another sample point. By assumption w.h.p. after $O(1)$ samples we will find a sample, which we call again x_0 , for which the dimension estimator succeeds. We have therefore found d_0 (say), and from the ball centered at x_0 of radius equal to the largest “good scale”, we estimate π_0 . At this point we assign to π_0 all the points that are no farther than δ from π_0 (we may choose $\delta = 0$ if there is no noise). We now repeat the above on the remaining points, till all the points have been assigned to some plane. After all the points have been assigned, a polishing step is performed: since we now have all the π_i ’s, we recompute the assignment of each point to the nearest π_i . Notice that the algorithm will succeed with high probability, as soon as each π_i has a fraction of points far enough from the intersection with other π_j ’s for which the dimension estimator is going to succeed. Recall that this estimator only requires $\mathcal{O}(d_i \log d_i)$ points in order to assess the dimension. The only remaining issue is the selection of δ , which we perform by estimating the noise variance empirically, as already done in the intrinsic dimension estimator. In fact, an even more robust procedure may be used, based on denoising and clustering the matrix of affinities between points and estimated candidates for the π_i ’s: such a procedure determines K and the assignment of the points to the π_i ’s at the same time: the details may be found in [24]. Finally, assuming pre-computation of any structure needed to compute nearest neighbors, and assuming that nearest neighbors may be computed in $O(\log n)$, the computational cost of the algorithm to find the pieces of planes and assign points to them is $\mathcal{O}(n \log n K \max_i d_i^2)$. If the assignment of points to planes is not required, and only the pieces of planes are requested, then randomization allows one to reduce the computational cost to $\mathcal{O}(K^2 \max_i d_i^2)$.

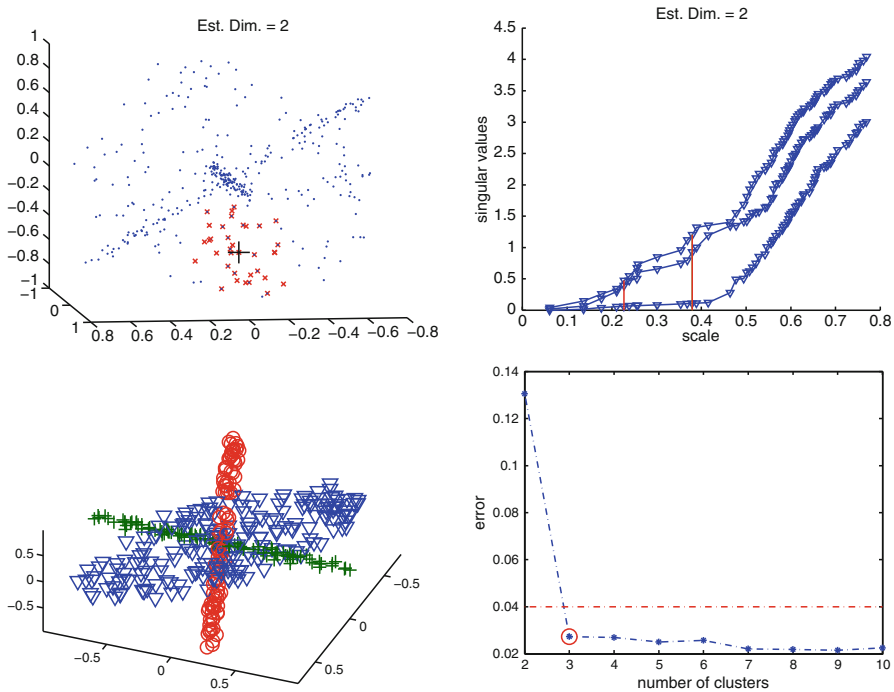


Fig. 10.15 Global spectral analysis of local subspaces for modeling and clustering hybrid linear data [24]. In this example, the given data set consists of points sampled around 2 lines and 1 plane in \mathbb{R}^3 . We first apply the intrinsic dimension estimator to the data and find a collection of good local scales with dimension estimates at many locations (see one such location (*black plus symbol*) in the top left figure, with region indicated by the red points) by analyzing the local singular values (see *top right*, the region between red vertical lines is the range of good scales). We then perform spectral analysis to integrate the local planes so that we recover the model as well as the underlying clusters (see *bottom left*; the misclassification rate is 2.5%). All we need so far, including for finding the intrinsic dimensions of the underlying planes, is the coordinates of the data plus knowledge of the number of clusters. When the number of clusters is not known, we detect it by looking at the errors associated with different choices (see *bottom right*; the red line indicates the model error which we can estimate by multiscale SVD)

We consider a simple example in Fig. 10.15, which in fact uses a more robust, less greedy version of the algorithm just described [24].

References

1. M. Aharon, M. Elad, and A. Bruckstein. K-SVD: Design of dictionaries for sparse representation. In *PROCEEDINGS OF SPARS 05'*, pages 9–12, 2005.
2. W.K. Allard, G. Chen, and M. Maggioni. Multiscale geometric methods for data sets II: Geometric wavelets. *in preparation*, 2010.

3. M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems 14 (NIPS 2001)*, pages 585–591. MIT Press, Cambridge, 2001.
4. M. Belkin and P. Niyogi. Using manifold structure for partially labelled classification. *Advances in NIPS*, 15, 2003.
5. M. Belkin and P. Niyogi. Semi-supervised learning on Riemannian manifolds. *Machine Learning*, 56(Invited Special Issue on Clustering):209–239, 2004. TR-2001-30, Univ. Chicago, CS Dept., 2001.
6. M. Belkin, P. Niyogi, and V. Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, (7):2399–2434, Nov. 2006.
7. J.J. Benedetto and M.W. Frazier eds. *Wavelets, Mathematics and Applications*. CRC Press, 1993.
8. P. Binev, A. Cohen, W. Dahmen, R.A. DeVore, and V. Temlyakov. Universal algorithms for learning theory part i: piecewise constant functions. *J. Mach. Learn. Res.*, 6:1297–1321, 2005.
9. P. Binev, A. Cohen, W. Dahmen, R.A. DeVore, and V. Temlyakov. Universal algorithms for learning theory part ii: piecewise polynomial functions. *Constr. Approx.*, 26(2):127–152, 2007.
10. I. Borg and P. Groenen. *Modern Multidimensional Scaling : Theory and Applications*. Springer, 1996.
11. S. Borovkova, R. Burton, and H. Dehling. Consistency of the Takens estimator for the correlation dimension. *Ann. Appl. Probab.*, 9(2):376–390, 1999.
12. J. Bourgain. On Lipschitz embedding of finite metric spaces into Hilbert space. *Isr. Journ. Math.*, pages 46–52, 1985.
13. F. Camastra and A. Vinciarelli. Intrinsic dimension estimation of data: An approach based on grassberger-procaccia's algorithm. *Neural Processing Letters*, 14(1):27–34, 2001.
14. F. Camastra and A. Vinciarelli. Estimating the intrinsic dimension of data with a fractal-based method. *IEEE P.A.M.I.*, 24(10):1404–1410, 2002.
15. Wenbo Cao and Robert Haralick. Nonlinear manifold clustering by dimensionality. *icpr*, 1:920–924, 2006.
16. K. Carter, A. O. Hero, and R. Raich. De-biasing for intrinsic dimension estimation. *Statistical Signal Processing, 2007. SSP '07. IEEE/SP 14th Workshop on*, pages 601–605, Aug. 2007.
17. Kevin M. Carter, Alfred O. Hero, and Raviv Raich. De-biasing for intrinsic dimension estimation. *Statistical Signal Processing, 2007. SSP '07. IEEE/SP 14th Workshop on*, pages 601–605, Aug. 2007.
18. K.M. Carter and A.O. Hero. Variance reduction with neighborhood smoothing for local intrinsic dimension estimation. *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 3917–3920, 31 2008-April 4 2008.
19. K.M. Carter and A.O. Hero. Variance reduction with neighborhood smoothing for local intrinsic dimension estimation. *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 3917–3920, 31 2008-April 4 2008.
20. Tony F. Chan and Jianhong Shen. *Image processing and analysis*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2005. Variational, PDE, wavelet, and stochastic methods.
21. M. Chaplain, M. Ganesh, and I. Graham. Spatio-temporal pattern formation on spherical surfaces: numerical simulation and application to solid tumor growth. *J. Math. Biology*, 42:387–423, 2001.
22. G. Chen and G. Lerman. Foundations of a multi-way spectral clustering framework for hybrid linear modeling. *Found. Comput. Math.*, 9:517–558, 2009. DOI 10.1007/s10208-009-9043-7.
23. G. Chen and G. Lerman. Spectral curvature clustering (sc). *Int. J. Comput. Vis.*, 81:317–330, 2009. DOI 10.1007/s11263-008-0178-9.
24. G. Chen and M. Maggioni. Multiscale geometric methods for data sets III: multiple planes. *in preparation*, 2010.
25. G. Chen and M. Maggioni. Multiscale geometric wavelets for the analysis of point clouds. *to appear in Proc. CISS 2010*, 2010.

26. M. Chen, J. Silva, J. Paisley, C. Wang, D. Dunson, and L. Carin. Compressive sensing on manifolds using a nonparametric mixture of factor analyzers: Algorithm and performance bounds. *IEEE Trans. Signal Processing*, 58(12):6140–6155, Dec. 2010.
27. Scott Shaobing Chen, David L. Donoho, and Michael A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.
28. C.K. Chui. *An introduction to wavelets*. Academic Press, San Diego, 1992.
29. R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS*, 102(21):7426–7431, 2005.
30. R.R. Coifman and S. Lafon. Diffusion maps. *Appl. Comp. Harm. Anal.*, 21(1):5–30, 2006.
31. J. Costa and A.O. Hero. Learning intrinsic dimension and intrinsic entropy of high dimensional datasets. In *Proc. of EUSIPCO*, Vienna, 2004.
32. J.A. Costa and A.O. Hero. Geodesic entropic graphs for dimension and entropy estimation in manifold learning. *Signal Processing, IEEE Transactions on*, 52(8):2210–2221, Aug. 2004.
33. J.A. Costa and A.O. Hero. Geodesic entropic graphs for dimension and entropy estimation in manifold learning. *Signal Processing, IEEE Transactions on*, 52(8):2210–2221, Aug. 2004.
34. I. Daubechies. *Ten lectures on wavelets*. Society for Industrial and Applied Mathematics, 1992.
35. G. David and S. Semmes. *Uniform Rectifiability and Quasiminimizing Sets of Arbitrary Codimension*. AMS.
36. G. David and S. Semmes. Singular integrals and rectifiable sets in \mathbf{R}^n : Au-delà des graphes lipschitziens. *Astérisque*, (193):152, 1991.
37. Guy David. <http://www.math.u-psud.fr/~gdavid/Notes-Parkcity.dvi>.
38. Guy David. Morceaux de graphes lipschitziens et intégrales singulières sur une surface. *Rev. Mat. Iberoamericana*, 4(1):73–114, 1988.
39. Guy David. *Wavelets and singular integrals on curves and surfaces*, volume 1465 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1991.
40. Guy David. *Wavelets and Singular Integrals on Curves and Surfaces*. Springer-Verlag, 1991.
41. Guy David and Stephen Semmes. *Analysis of and on uniformly rectifiable sets*, volume 38 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1993.
42. D. L. Donoho and Ana G Flesia. Can recent innovations in harmonic analysis ‘explain’ key findings in natural image statistics? *Network: Comput. Neural Syst.*, 12:371–393, 2001.
43. D. L. Donoho and C. Grimes. When does isomap recover natural parameterization of families of articulated images? Technical Report Tech. Rep. 2002-2027, Department of Statistics, Stanford University, August 2002.
44. D. L. Donoho and Carrie Grimes. Hessian eigenmaps: new locally linear embedding techniques for high-dimensional data. *Proc. Nat. Acad. Sciences*, pages 5591–5596, March 2003. also tech. report, Statistics Dept., Stanford University.
45. D. L. Donoho, O. Levi, J.-L. Starck, and V. J. Martinez. Multiscale geometric analysis for 3-d catalogues. Technical report, Stanford Univ., 2002.
46. A. M. Farahmand and C. Szepesv and J.-Y. Audibert. Manifold-adaptive dimension estimation. *Proc. I.C.M.L.*, 2007.
47. A. M. Farahmand, Cs. Szepesvári, and J.-Y. Audibert. Manifold-adaptive dimension estimation. In *Proceedings of the 24th international conference on Machine learning*, page 265–272, 2007.
48. S. Atev G. Chen and G. Lerman. Kernel spectral curvature clustering (kscc). In *The 4th ICCV International Workshop on Dynamical Vision*, Kyoto, Japan, 2009.
49. Peter Grassberger and Itamar Procaccia. Measuring the strangeness of strange attractors. *Phys. D*, 9(1-2):189–208, 1983.
50. G. Haro, G. Randall, and G. Sapiro. Translated poisson mixture model for stratification learning. *Int. J. Comput. Vision*, 80(3):358–374, 2008.
51. Gloria Haro, Gregory Randall, and Guillermo Sapiro. Translated poisson mixture model for stratification learning. *Int. J. Comput. Vision*, 80(3):358–374, 2008.
52. X. He, S. Yan, Y. Hu, P. Niyogi, and H.-J. Zhang. Face recognition using laplacianfaces. *IEEE Trans. pattern analysis and machine intelligence*, 27(3):328–340, 2005.

53. M. Hein and Y. Audibert. Intrinsic dimensionality estimation of submanifolds in euclidean space. In S. Wrobel De Raedt, L., editor, *ICML Bonn*, pages 289 – 296, 2005.
54. Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Ann. Stat.*, 29(2):295–327, April 2001.
55. Peter W. Jones. Rectifiable sets and the traveling salesman problem. *Invent. Math.*, 102(1):1–15, 1990.
56. P.W. Jones. Rectifiable sets and the traveling salesman problem. *Inventiones Mathematicae*, 102:1–15, 1990.
57. G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, 1999.
58. V. I. Koltchinskii. Empirical geometry of multivariate data: a deconvolution approach. *Ann. Stat.*, 28(2):591–629, 2000.
59. R. Krauthgamer, J. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics, 2004.
60. S. Lafon. *Diffusion maps and geometric harmonics*. PhD thesis, Yale University, 2004.
61. E. Levina and P. Bickel. Maximum likelihood estimation of intrinsic dimension. In *Advances in NIPS 17, Vancouver, Canada*, 2005.
62. Elizaveta Levina and Peter J. Bickel. Maximum likelihood estimation of intrinsic dimension. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 777–784. MIT Press, Cambridge, MA, 2005.
63. A.V. Little, Y.-M. Jung, and M. Maggioni. Multiscale estimation of intrinsic dimensionality of data sets. In *Proc. A.A.A.I.*, 2009.
64. A.V. Little, J. Lee, Y.-M. Jung, and M. Maggioni. Estimation of intrinsic dimensionality of samples from noisy low-dimensional manifolds in high dimensions with multiscale SVD. In *Proc. S.S.P.*, 2009.
65. A.V. Little, M. Maggioni, and L. Rosasco. Multiscale geometric methods for data sets I: Estimation of intrinsic dimension. *in preparation*, 2010.
66. P-C. Lo. Three dimensional filtering approach to brain potential mapping. *IEEE Tran. on bio-medical engineering*, 46(5):574–583, 1999.
67. S. Mahadevan, K. Ferguson, S. Osentoski, and M. Maggioni. Simultaneous learning of representation and control in continuous domains. In *AAAI*. AAAI Press, 2006.
68. S. Mahadevan and M. Maggioni. Value function approximation with diffusion wavelets and laplacian eigenfunctions. In *University of Massachusetts, Department of Computer Science Technical Report TR-2005-38; Proc. NIPS 2005*, 2005.
69. Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. In *ICML*, page 87, 2009.
70. S.G. Mallat. *A wavelet tour in signal processing*. Academic Press, 1998.
71. Benoit B. Mandelbrot and Richard L. Hudson. *The (mis)behavior of markets*. Basic Books, New York, 2004. A fractal view of risk, ruin, and reward.
72. A. Ng, M. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm, 2001.
73. P. Niyogi, I. Matveeva, and M. Belkin. Regression and regularization on large graphs. Technical report, University of Chicago, Nov. 2003.
74. Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17:1617–1642, 2007.
75. M. Raginsky and S. Lazebnik. Estimation of intrinsic dimensionality using high-rate vector quantization. *Proc. NIPS*, pages 1105–1112, 2005.
76. ST Roweis and LK Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
77. M. Rudelson. Random vectors in the isotropic position. *J. of Functional Analysis*, 164(1):60–72, 1999.
78. L.K. Saul, K.Q. Weinberger, F.H. Ham, F. Sha, and D.D. Lee. *Spectral methods for dimensionality reduction*, chapter Semisupervised Learning. MIT Press, 2006.
79. F. Sha and L.K. Saul. Analysis and extension of spectral methods for nonlinear dimensionality reduction. *Proc. ICML*, pages 785–792, 2005.

80. J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE PAMI*, 22:888–905, 2000.
81. Jack Silverstein. On the empirical distribution of eigenvalues of large dimensional information-plus-noise type matrices. *Journal of Multivariate Analysis*, 98:678–694, 2007.
82. A. Szlam and G. Sapiro. Discriminative k -metrics. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1009–1016, 2009.
83. A.D. Szlam, M. Maggioni, and R.R. Coifman. Regularization on graphs with function-adapted diffusion processes. *JMLR*, (9):1711–1739, Aug 2008.
84. A.D. Szlam, M. Maggioni, R.R. Coifman, and J.C. Bremer Jr. Diffusion-driven multiscale analysis on manifolds and graphs: top-down and bottom-up constructions. volume 5914-1, page 59141D. SPIE, 2005.
85. Floris Takens. On the numerical determination of the dimension of an attractor. In *Dynamical systems and bifurcations (Groningen, 1984)*, volume 1125 of *Lecture Notes in Math.*, pages 99–106. Springer, Berlin, 1985.
86. J.B. Tenenbaum, V. de Silva, and J.C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000.
87. M. B. Wakin, D. L. Donoho, H. Choi, and R. G. Baraniuk. The multiscale structure of non-differentiable image manifolds. In *SPIE Wavelets XI*, San Diego, July 2005.
88. K.Q. Weinberger, F. Sha, and L.K. Saul. Learning a kernel matrix for nonlinear dimensionality reduction. *Proc. ICML*, pages 839–846, 2004.
89. Mladen Victor Wickerhauser. *Adapted Wavelet Analysis from Theory to Software*. A K Peters Ltd., Wellesley, MA, 1994. With a separately available computer disk (IBM-PC or Macintosh).
90. Z. Zhang and H. Zha. Principal manifolds and nonlinear dimension reduction via local tangent space alignment. Technical Report CSE-02-019, Department of computer science and engineering, Pennsylvania State University, 2002.
91. M. Zhou, H. Chen, J. Paisley, L. Ren, G. Sapiro, and L. Carin. Non-parametric Bayesian dictionary learning for sparse image representations. In *Neural and Information Processing Systems (NIPS)*, 2009.
92. Xiaojin Zhu, Zoubin Ghahramani, and John Lafferty. Semi-supervised learning using gaussian fields and harmonic functions. In *ICML*, pages 912–919, 2003.

Chapter 11

Signal Ensemble Classification Using Low-Dimensional Embeddings and Earth Mover's Distance

Linh Lieu and Naoki Saito

Abstract Instead of classifying individual signals, we address classification of objects characterized by signal ensembles (i.e., collections of signals). Such necessity arises frequently in real situations: e.g., classification of video clips or object classification using acoustic scattering experiments to name a few. In particular, we propose an algorithm for classifying signal ensembles by bringing together well-known techniques from various disciplines in a novel way. Our algorithm first performs the dimensionality reduction on training ensembles using either the linear embeddings (e.g., Principal Component Analysis (PCA), Multidimensional Scaling (MDS)) or the nonlinear embeddings (e.g., the Laplacian eigenmap (LE), the diffusion map (DM)). After embedding training ensembles into a lower-dimensional space, our algorithm extends a given test ensemble into the trained embedding space, and then measures the “distance” between the test ensemble and each training ensemble in that space, and classify it using the nearest neighbor method. It turns out that the choice of this ensemble distance measure is critical, and our algorithm adopts the so-called Earth Mover's Distance (EMD), a robust distance measure successfully used in image retrieval and image registration. We will demonstrate the performance of our algorithm using two real examples: classification of underwater objects using multiple sonar waveforms; and classification of video clips of digit-speaking lips. This article also provides a concise review on the several key concepts in statistical learning such as PCA, MDS, LE, DM, and EMD as well as the practical issues including how to tune parameters, which will be useful for the readers interested in numerical experiments.

N. Saito (✉)

Department of Mathematics, University of California, One Shields Avenue, Davis,
CA 95616, USA

e-mail: saito@math.ucdavis.edu

11.1 Introduction

Many problems in pattern recognition often require comparison between ensembles of signals (i.e., points in a high-dimensional space) instead of comparing individual signals. For example, in a visual speech recognition problem, two or more clips of recorded video are compared for similar patterns. Typically, a video clip consists of a sequence of images called video frames. We can view a video frame as a point and a video clip as an ensemble of points in the high-dimensional image space. Thus, comparing two video clips is equivalent to comparing two ensembles of points in the image space. As a second example, consider the task of identifying an object on the ocean floor given a set of sonar waveforms reflected from the object. This requires comparing the ensembles of waveforms reflected from the unknown object with those reflected from known objects. Then the unknown object is identified when a match is made.

In this article, we propose an algorithm for classifying such signal ensembles. More specifically, let $\mathcal{X} := \bigcup_{i=1}^M \mathcal{X}^i \subset \mathbb{R}^d$ be a collection of M training *ensembles* $\mathcal{X}^i := \{\mathbf{x}_1^i, \dots, \mathbf{x}_{m_i}^i\}$, $i = 1, \dots, M$ where $m_i \in \mathbb{N}$ is the number of signals contained in the ensemble \mathcal{X}^i and $\mathbf{x}_k^i \in \mathbb{R}^d$ is an individual signal (a vector of length d). We assume that each training ensemble \mathcal{X}^i has a unique label (or class name) among C possible labels. Let $\mathcal{Y} := \bigcup_{j=1}^N \mathcal{Y}^j \subset \mathbb{R}^d$ be a collection of test (i.e., unlabeled) ensembles where $\mathcal{Y}^j := \{\mathbf{y}_1^j, \dots, \mathbf{y}_{n_j}^j\}$ is the j th test ensemble consisting of n_j signals of length d . Our goal is to classify each \mathcal{Y}^j , $j = 1, \dots, N$, to one of the possible C classes given the training ensembles \mathcal{X} . Note that this task is different from classifying each signal $\mathbf{y}_k^j \in \mathcal{Y}^j$ individually.

Our proposed algorithm consists of two main stages. The first (or training) stage performs the dimensionality reduction without losing important features of the training ensembles followed by constructing a compact representation called a *signature* – a generalized version of a histogram – that well represents an essence of the ensemble in the reduced space. The second (or classification) stage embeds a given test ensemble into the reduced space obtained in the first stage followed by classifying it to the label of the training ensemble whose signature is most similar to that of the test ensemble. Here, how to define the similarity or the distance measure between signatures will be the key issue as we will discuss below.

11.1.1 Dimensionality Reduction

Let us now discuss the first stage in more details. The dimensionality reduction of the input data is necessary because modern technologies often generate data of extremely high dimension. For example, a small 128×128 gray-scale image has dimension $d = 16384$. Such high dimensionality makes data analysis inefficient and inaccurate. Fortunately, the data that we encounter often have low intrinsic

dimensionality. For example, consider the set of all 128×128 gray-scale images taken of an object under fixed lighting by a moving camera. This is a subset of \mathbb{R}^{16384} possessing the natural structure of a low-dimensional manifold with its dimensionality defined by the degrees of freedom of the camera [3]. In short, it is desirable and in fact indispensable to find an effective low-dimensional representation of the data for the task at hand.

Perhaps, the most well-known techniques are Principal Component Analysis (PCA) also known as Karhunen–Loève Transform (KLT) [27, Chap. 8], [17, Sect. 14.5], and Multidimensional Scaling (MDS) [7], [27, Chap. 14], [17, Sect. 14.8]. Despite their shortcomings (e.g., inability to capture *local* geometric information), they have been the mainstay of dimensionality reduction. They have been applied to a variety of applications including human face characterization [18, 35], pose estimation of 3D objects from image sequences [28], to name only a few; see also [17, Sect. 14.5] for more examples.

Many nonlinear methods for dimensionality reduction have been proposed in order to improve the shortcomings of PCA/MDS; see e.g., [12, 32, 41] for some specific proposals and [24] for a comparative review on many such techniques. Unlike classical methods such as PCA and MDS, nonlinear methods in general offer the advantage of preserving local geometry while achieving dimensionality reduction. Providing the example of a moving camera in the paragraph above for motivation, Belkin and Niyogi [2–4] were among the group of scientists who proposed a nonlinear dimensionality reduction algorithm that explicitly considers the manifold structure that may very well be the intrinsic geometry of the data. They proposed the *Laplacian eigenmap* (LE): an embedding of data into a reduced space spanned by the eigenvectors of the graph Laplacian defined on a similarity graph constructed from the training data while preserving the local geometry of the data. If data arise from non-uniform sampling of a manifold, however, embedding via a Laplacian eigenmap may result in distortion of the manifold when embedded into the reduced space (see [21] for some examples). Sensitivity to sampling densities may be a serious drawback in certain cases. For this reason, Coifman and Lafon [9, 21] proposed a density-invariant normalization of the weights on the graph before computing the graph Laplacian. This would eliminate sensitivity to sampling densities of the Laplacian eigenmaps. Furthermore, the authors defined *diffusion maps* from the eigenvalues and eigenvectors of the diffusion operator (defined in (11.14) below) and provided an intuitive interpretation of how point clustering in a diffusion coordinate system is linked to a Markov random walk on the weighted graph; see e.g., [20] for the details.

In this chapter, we will particularly focus on PCA, MDS, Laplacian eigenmaps, and diffusion maps as our dimensionality reduction methods and compare their performance.

11.1.2 *Nearest Neighbor Classification and Its Base Distance Measure*

The final step of the second stage of our proposed algorithm is to actually classify each test ensemble into one of the possible classes by comparing its signature with those of the training ensembles computed in the first stage. The label of the training ensemble whose signature is “closest” to that of a given test ensemble is assigned to it. In other words, our algorithm adopts the so-called k -nearest neighbor (k -NN) classifier with $k = 1$ on the signatures in the reduced embedding space. The k -NN classifier has been quite successful in many classification problems despite its simplicity; see e.g., [17, Sect. 13.3] for the details. Although our algorithm can easily accommodate a different value of $k > 1$, we will stick with $k = 1$ in this article mainly because: (1) choosing the optimal k requires a relatively large number of training ensembles and (2) a famous result of Cover and Hart [11] claims that the error rate of the 1-NN classifier is never more than twice the Bayes (i.e., optimal) rate asymptotically.

The more important issue in this second stage is the choice of the base similarity or distance measure for the nearest neighbor classifier. There are many such measures for comparing given two signatures: e.g., the usual Euclidean distance, the Hellinger distance, the Kullback–Leibler divergence, etc.; see [1] for the review of many of these distances. The most robust one that is also most suitable for our problem turns out to be the so-called *Earth Mover’s Distance* (EMD), which has been successfully applied to many practical problems including image retrieval from a large image database [33, 34], image registration and warping [16], to name a few. To the best of our knowledge, however, using EMD for classification of objects characterized by signal ensembles has not been addressed by others. In our case, each signature may consist of different number of clusters or “bins” in the reduced embedding space. Hence, the conventional distances measure such as the Euclidean distance, the Hellinger distance, etc., are not suitable for our classification problem.

11.1.3 *Related Works*

In this section, we will briefly describe two works most related to ours. In [42], the authors presented a method to characterize signal ensembles using various probabilistic distance measures in the reproducing kernel Hilbert space (RKHS). First, the signal ensembles are mapped to an RKHS via a nonlinear map using either a polynomial kernel or the Gaussian kernel. Then the probabilistic distance measures (e.g., the Bhattacharyya distance, the Kullback–Leibler divergence, the Chernoff distance, etc.) are computed in the RKHS space to determine the similarities between signal ensembles. Our proposed method and [42] share the common objective of classifying signal ensembles. Other than using a nonlinear mapping to transform the data from its original space, however, our algorithm and [42] are two completely different

approaches. For example, the approach of [42] assumes that the data after mapped to an RKHS obey the normal (or Gaussian) distribution, and they proceed to estimate the parameters (the mean vector and the covariance matrix) of the normal distribution. Such an assumption is too strong and cannot be guaranteed in practice in our opinion even if one finds a good nonlinear kernel function.

In [19], Lafon, Keller, and Coifman utilized diffusion maps and the out-of-sample extension scheme for signal ensemble matching problems. Our proposed algorithm owes them quite a bit: we too use diffusion maps and their out-of-sample extension scheme, which will be reviewed in Sect. 11.2.4 in detail. An important difference between our algorithm and theirs, however, is the use of EMD as the similarity/distance measure instead of the Hausdorff distance (HD) they adopted; see (11.22) for the definition of HD and Sect. 11.3 for the detailed discussion. We also note that the techniques we employ in the signature construction stage for EMD-based classification is not limited to the framework of the diffusion maps; in fact, we also use the Laplacian eigenmaps and PCA in our numerical experiments below. Furthermore, the advantage of EMD over HD is its robustness to outliers. As we will see in Sect. 11.5 below, HD is very sensitive to outliers. In contrast, the EMD between the signatures measures the differences between the distributions of points within the reduced space. Therefore it is less affected by outliers, which leads to more reliable results.

11.1.4 Article Organization

In Sect. 11.2, we will review the basics of the particular low-dimensional embedding techniques for dimensionality reduction, i.e., Multidimensional Scaling/Principal Component Analysis, the Laplacian eigenmaps, and the diffusion maps. Section 11.3 reviews the key player in our proposed algorithm: the Earth Mover's Distance. Section 11.4 describes our proposed algorithm for signal ensemble classification in detail including the practical issues such as the tuning parameter selection that are important to successfully embed high-dimensional signals into the lower dimensional spaces using the diffusion maps/Laplacian eigenmaps. Section 11.5 describes the results of our numerical experiments on two real datasets: underwater object classification using acoustic signals and classification of video clips of digit-speaking lips. Finally, we will conclude in Sect. 11.6 with discussion.

11.2 Dimensionality Reduction/Low-Dimensional Embeddings

In this section, we will review the embeddings of high-dimensional data into a low-dimensional space using MDS/PCA, Laplacian eigenmaps, and diffusion maps and discuss their properties that allow us to achieve meaningful dimensionality reduction

as well as their similarity and difference. We will also review an algorithm proposed in [19] for extending these embedding maps from training data to test data.

First of all, let $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{m_*}\} \subset \mathbb{R}^d$ be a set of the available training data points (or signals). Note that using the notations defined in Introduction, the total number of training data points $m_* := \sum_{i=1}^M m_i$. In this article, we assume the finite number of data points both for training and test signals. Since $\mathcal{X} \subset \mathbb{R}^d$, we can define a natural dissimilarity measure δ that gives a sense of affinity between any two data points in \mathcal{X} . The usual Euclidean distance is most commonly used as δ although there are many other possibilities. Let X be a data matrix of size $d \times m_*$ (i.e., each column vector is a training signal)¹, and let $\Psi : \mathcal{X} \subset \mathbb{R}^d \rightarrow \mathbb{R}^s$ be an embedding map from high to low dimensional spaces where $1 \leq s \ll \min(d, m_*)$ is the reduced dimension. With a slight abuse of notation, let $\Psi(X)$ denote a matrix of embedded points of \mathcal{X} in \mathbb{R}^s , i.e.,

$$\Psi(X) := (\Psi(\mathbf{x}_1), \dots, \Psi(\mathbf{x}_{m_*})) = \begin{bmatrix} \Psi_1^T \\ \vdots \\ \Psi_s^T \end{bmatrix} = \begin{bmatrix} \psi_1(\mathbf{x}_1) & \dots & \psi_1(\mathbf{x}_{m_*}) \\ \vdots & & \vdots \\ \psi_s(\mathbf{x}_1) & \dots & \psi_s(\mathbf{x}_{m_*}) \end{bmatrix} \in \mathbb{R}^{s \times m_*}. \quad (11.1)$$

11.2.1 Multidimensional Scaling and Principal Component Analysis

Multidimensional Scaling (MDS) [7], [27, Chap. 14], [17, Sect. 14.8] seeks a low-dimensional embedding that minimizes the following quantity:

$$J_{\text{MDS}}(\Psi) := \sum_{i=1}^{m_*} \sum_{j=i}^{m_*} (\delta(\mathbf{x}_i, \mathbf{x}_j) - \delta(\Psi(\mathbf{x}_i), \Psi(\mathbf{x}_j)))^2, \quad (11.2)$$

where $\delta(\cdot, \cdot)$ is the “distance” between the two points in the appropriate Euclidean space. In other words, MDS seeks a low-dimensional representation of the input data that preserves the pairwise distances as well as possible. Although it is possible to use any distances measure (i.e., metric) as δ , the most common choice by far is the Euclidean distance, e.g., $\delta(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|$ where $\|\cdot\|$ is the ℓ^2 -norm. We note that the embedding minimizing J_{MDS} could be nonlinear in general.

Instead of using the distance or dissimilarity, it is also possible to use the similarity among the data points. If the similarity is defined by the centered correlation as

$$\alpha(\mathbf{x}_i, \mathbf{x}_j) := (\mathbf{x}_i - \bar{\mathbf{x}})^T (\mathbf{x}_j - \bar{\mathbf{x}}), \quad \bar{\mathbf{x}} := \frac{1}{m_*} \sum_{i=1}^{m_*} \mathbf{x}_i, \quad (11.3)$$

¹ In the statistics literature, the data matrix is often defined as the transpose of our definition, i.e., each row vector is a training signal.

then instead of (11.2) one can consider the embedding minimizing the following

$$J_{\text{CS}}(\Psi) := \sum_{i=1}^{m_*} \sum_{j=i}^{m_*} (\alpha(\mathbf{x}_i, \mathbf{x}_j) - \alpha(\Psi(\mathbf{x}_i), \Psi(\mathbf{x}_j)))^2. \quad (11.4)$$

This is called *Classical MDS* or *Classical Scaling* (CS). Using the matrix-vector notation, and let \tilde{X} be the centered data matrix, i.e., the mean of the column vectors of X is subtracted from each column of X , which can be written as $\tilde{X} = XH$ where $H := I - \mathbf{1}\mathbf{1}^T/m_*$ is the centering matrix of size $m_* \times m_*$ and $\mathbf{1} := (1, \dots, 1)^T \in \mathbb{R}^{m_*}$. Let $\tilde{\Psi}(X) := \Psi(X)H$. (Note that if we assume Ψ is a linear map, then $\tilde{\Psi}(X) = \Psi(\tilde{X})$, but let us proceed without assuming this linearity now.) Then, (11.4) can be written as

$$J_{\text{CS}}(\Psi) = \left\| \tilde{X}^T \tilde{X} - \tilde{\Psi}(X)^T \tilde{\Psi}(X) \right\|_F^2, \quad (11.5)$$

where $\|\cdot\|_F$ is the Frobenius norm. Since $\text{rank}(\Psi(X)) = s$ and $\text{rank}(\tilde{\Psi}(X)) = s - 1$, the minimization of $J_{\text{CS}}(\Psi)$ implies that $\tilde{\Psi}(X)^T \tilde{\Psi}(X)$ should be the best $\text{rank}(s - 1)$ -approximation of $\tilde{X}^T \tilde{X}$. Let $\tilde{X} = U \Sigma V^T$ be the singular value decomposition of \tilde{X} . Then, $\tilde{X}^T \tilde{X} = V \Sigma^T \Sigma V^T$. Hence, the best $\text{rank}-(s - 1)$ approximation of $\tilde{X}^T \tilde{X}$ is $\sum_{k=1}^{s-1} \sigma_k \mathbf{v}_k \mathbf{v}_k^T$ where σ_k are sorted in the non-increasing order. This implies that

$$\tilde{\Psi}(X) = \begin{bmatrix} \sigma_1 \mathbf{v}_1^T \\ \vdots \\ \sigma_{s-1} \mathbf{v}_{s-1}^T \end{bmatrix} = \Sigma_{s-1} V_{s-1}^T,$$

where $\Sigma_{s-1} = \text{diag}(\sigma_1, \dots, \sigma_{s-1})$ and $V_{s-1} = (\mathbf{v}_1, \dots, \mathbf{v}_{s-1})$. However, because $\mathbf{1} \in \mathbb{R}^{m_*}$ is an eigenvector of \tilde{X} corresponding to the eigenvalue 0 (this can be easily seen from $H\mathbf{1} = \mathbf{0}$), the mean of the column vectors of the matrix $\sum_k V_k^T$ is $\mathbf{0}$ as long as $\sigma_k \geq 0$ i.e., $k \leq \text{rank}(\tilde{X})$ because $V_k^T \mathbf{1} = \mathbf{0}$. In other words, $\sum_k V_k^T$ is already centered for each such k . This also implies that we can set $\Psi(X) = \sum_s V_s^T$ as our s -dimensional embedding as long as $1 \leq s \leq \text{rank}(\tilde{X})$, which is usually the case since $s \ll \min(d, m_*)$ and the training signals normally linearly independent, i.e., $\text{rank}(\tilde{X}) = \min(d, m_*) - 1$.

Now it is clear that the classical scaling is equivalent to the popular Principal Component Analysis (PCA), which was pointed out originally by [15]; see also [27, Sect. 14.3], [7, Sect. 24.1]. Since the first s PCA basis vectors are the eigenvectors of the sample covariance matrix $\tilde{X} \tilde{X}^T / m_*$ corresponding to the largest s eigenvalues, these are the first s column vectors of U , which form U_s . From the SVD of \tilde{X} , we have $\tilde{X} V_s = U_s \Sigma_s$ or equivalently $U_s = \tilde{X} V_s \Sigma_s^{-1}$. Now, the first s principal component of the centered data matrix \tilde{X} is simply the expansion of \tilde{X} with respect to U_s . Hence, those principal components (or the expansion coefficients) are

$$U_s^T \tilde{X} = \Sigma_s^{-1} V_s^T \tilde{X}^T \tilde{X} = [\Sigma_s | O_{s \times (m_* - s)}] V^T = \Sigma_s V_s^T.$$

In other words, this is exactly the same as CS. Hence, we use the term PCA interchangeably with CS or classical MDS in this article, and define

$$\Psi_{\text{PCA}}(X) := \Sigma_s V_s^T, \quad XH = U \Sigma V^T. \quad (11.6)$$

One of the major drawbacks of PCA and CS is their inability to incorporate *local* geometric information. This is clear from (11.3), which gives a small value even if \mathbf{x}_i and \mathbf{x}_j are very close as long as $\mathbf{x}_i - \bar{\mathbf{x}}$ and $\mathbf{x}_j - \bar{\mathbf{x}}$ are almost perpendicular.

11.2.2 Laplacian Eigenmaps

Unlike PCA/CS, Laplacian eigenmaps [2, 3] incorporate *local* geometric information in the ambient space \mathbb{R}^d when constructing a low-dimensional embedding. To explain this more precisely, let us define the weight function $w(\mathbf{x}_i, \mathbf{x}_j)$. A canonical example of this weight function is the Gaussian:

$$w_\varepsilon(\mathbf{x}_i, \mathbf{x}_j) := e^{-(\delta(\mathbf{x}_i, \mathbf{x}_j)/\varepsilon)^2}, \quad (11.7)$$

where $\varepsilon > 0$ is the scale parameter and $\delta(\cdot, \cdot)$ is an appropriate metric in the original ambient space \mathbb{R}^d . Although there are other choices, we will use the Euclidean distance (i.e., ℓ^2 -norm) as δ throughout this article. Such a weight function w_ε defines the notion of a local neighborhood at each point $\mathbf{x} \in \mathcal{X}$ via the affinity between \mathbf{x} and other points in \mathcal{X} , and the value of the scale parameter ε specifies the size of this neighborhood. Moreover, as explained in [3], when the dataset \mathcal{X} approximately lies on a lower-dimensional manifold in \mathbb{R}^d , using the weights w_ε on the graph constructed from the data points corresponds to an approximation of the heat kernel on this manifold. We also note that choosing the appropriate value of ε for a given training dataset \mathcal{X} is quite important yet difficult in practice. We will discuss more about it in Sect. 11.4.2.

Now, we seek an embedding map Ψ that minimizes the following objective function

$$J_{\text{LE}}(\Psi) := \sum_{i=1}^{m_*} \sum_{j=1}^{m_*} \|\Psi(\mathbf{x}_i) - \Psi(\mathbf{x}_j)\|^2 w_\varepsilon(\mathbf{x}_i, \mathbf{x}_j). \quad (11.8)$$

Using the matrix-vector notation, this leads to the following optimization problem:

$$\min_{\Psi(X) \in \mathbb{R}^s \times m_*} \text{tr}(\Psi(X) L \Psi(X)^T) \quad \text{subject to } \Psi(X) D \Psi(X)^T = I, \quad (11.9)$$

where the matrices are defined as

$$W := (w_\varepsilon(\mathbf{x}_i, \mathbf{x}_j)), \quad D := \text{diag} \left(\sum_j w_\varepsilon(\mathbf{x}_1, \mathbf{x}_j), \dots, \sum_j w_\varepsilon(\mathbf{x}_{m_*}, \mathbf{x}_j) \right).$$

The matrices of size $m_* \times m_*$, $L := D - W$, D , and W are called the (unnormalized) *graph Laplacian*, the degree matrix (diagonal), and the weighted adjacency matrix, respectively. The constraint $\Psi(X)D\Psi(X)^T = I$ is imposed to prevent Ψ from mapping X to a subspace whose dimension is lower than s . It is now straightforward to show that the minimizer $\Psi(X)$ of (11.9) satisfies the following generalized eigenvalue problem:

$$L\Psi(X)^T = D\Psi(X)^T\Lambda,$$

where $\Lambda = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{s-1})$ is a $s \times s$ diagonal matrix listing the eigenvalues in the *non-decreasing* order, i.e., $0 \leq \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_s$. The above generalized eigenvalue problem is equivalent to

$$L_{\text{rw}}\Psi_{\text{rw}}(X)^T = \Psi_{\text{rw}}(X)^T\Lambda_{\text{rw}}, \quad L_{\text{rw}} := D^{-1}L = I - D^{-1}W. \quad (11.10)$$

Then, the *Laplacian eigenmap* of the training dataset \mathcal{X} is defined as $\Psi_{\text{rw}}(X) \in \mathbb{R}^{s \times m_*}$.

Four important remarks are in order. First, notice that the row vectors of $\Psi_{\text{rw}}(X)$ – if they are transposed – form the eigenvectors of L_{rw} . Second, note that $\lambda_0 = 0$ and the corresponding eigenvector is $\mathbf{1} \in \mathbb{R}^{m_*}$, i.e., $L_{\text{rw}}\mathbf{1} = \mathbf{0}$, because if we view the data points as nodes and w_e as the edge weights of a graph, then our graph is *fully connected*. This constant vector $\mathbf{1}$ does not provide any information (except that the graph is fully connected, which we know a priori) and is useless for the embedding. Hence, we solve (11.10) for the $(s+1)$ -dimensional embedding $\Psi(X)$, and use its 2nd to the last rows as the s -dimensional embedding $\Psi_{\text{rw}}(X)$. Finally, $D^{-1}W$ is now a row-stochastic matrix (i.e., the sum of each row is 1, representing a transition probability distribution of a random walk over the graph). Hence, this multiplication of D^{-1} to L or W from left is called the *row-stochastic normalization*. That is the reason why we used the subscript *rw* in (11.10), which stands for ‘random walk’.

Second, instead of (11.10), one can consider a different normalization of the graph Laplacian, which is sometimes useful for numerical computation and studying its relationship to the spectral geometry (see, e.g., [8]). This is based on the following eigenvalue problem

$$L_{\text{sym}}\Psi_{\text{sym}}(X)^T = \Psi(X)^T_{\text{sym}}\Lambda_{\text{sym}}, \quad L_{\text{sym}} := D^{-\frac{1}{2}}LD^{-\frac{1}{2}} = I - D^{-\frac{1}{2}}WD^{-\frac{1}{2}}. \quad (11.11)$$

Both L_{rw} and L_{sym} are called the *normalized* graph Laplacians. These two are related in the following manner:

$$\Psi_{\text{rw}}(X) = \Psi_{\text{sym}}(X)D^{-\frac{1}{2}}, \quad \Lambda_{\text{rw}} = \Lambda_{\text{sym}}.$$

For more about the differences between these two normalizations, see the excellent tutorial paper on the spectral clustering [23].

Third, when the data arise from non-uniform sampling of a manifold, embedding via a Laplacian eigenmap may result in distortion of the manifold when embedded

into the reduced space. Such sensitivity to sampling densities may be a serious drawback in certain cases; see [21] for some examples. In fact, this was one of the motivations for Coifman and Lafon [9, 21] to develop a density-invariant normalization of the weights on the graph before computing the graph Laplacian, which we will explain in the next section. This normalization would eliminate sensitivity to sampling densities of the Laplacian eigenmaps. However, this normalization was not presented in the original definition of Laplacian eigenmaps [3] and is one of two important distinctions between the diffusion maps and the Laplacian eigenmaps. (The other is the incorporation of the eigenvalues into the map). Therefore, when we refer to Laplacian eigenmaps in this article, we stand by its original definition. When comparing the performance of different methods in our numerical experiments, we perform density-invariant normalization only when computing diffusion maps.

Fourth, the eigenvalues of the graph Laplacians have been studied extensively for many years; e.g., [8] is a fundamental reference, whose definition of the graph Laplacian is in fact L_{sym} . They have also been applied for a variety of purposes including our own work on characterizing neuronal dendrite patterns [38]. The eigenvectors of the graph Laplacians have also drawn attention recently from computer science and discrete mathematics communities. See e.g., [6] for the basic reference, and [23] for the application to spectral clustering.

11.2.3 Diffusion Maps

Following the work of Coifman and Lafon [9, 21], we construct the diffusion operator on \mathcal{X} as follows. First, we apply the *density-invariant normalization* to the weighted adjacency matrix W , which can be written as

$$\tilde{W} := D^{-1}WD^{-1}. \quad (11.12)$$

Then, we apply the row-stochastic normalization as the Laplacian eigenmap algorithm, i.e.,

$$\tilde{A}_{\text{rw}} := \tilde{D}^{-1}\tilde{W}, \quad (11.13)$$

where \tilde{D} is the degree matrix (diagonal) of \tilde{W} . Instead of the graph Laplacian L_{rw} used in the Laplacian eigenmaps, the diffusion map algorithm focuses on the normalized weighted adjacency matrix \tilde{A}_{rw} . One can interpret \tilde{A}_{rw} as the *transition matrix* of a random walk on \mathcal{X} or the *diffusion operator* on \mathcal{X} . Then, the t th power \tilde{A}_{rw}^t corresponds to running such random walk t steps forward in time. The information (i.e., a random walk) propagates more easily and quickly among the regions of high affinity (i.e., data points connected with the larger values of w_ϵ) than those of low affinity. This is essentially how we capture the local geometric information in the data. The diffusion map algorithm then performs the eigenanalysis

$$\tilde{A}_{\text{rw}} \Psi_{\text{DM}}(X)^T = \Psi_{\text{DM}}(X)^T \Lambda_{\text{DM}}, \quad (11.14)$$

where the eigenvalues are sorted in *non-increasing* order. The *diffusion map* is then defined (with a slight abuse of notation) as

$$\Psi_{\text{DM}}^t(X) := \Lambda_{\text{DM}}^t \Psi_{\text{DM}}(X). \quad (11.15)$$

It is important to point out its relationship with the Laplacian eigenmap (if \tilde{W} is used for constructing L_{rw} instead of W in (11.10)):

$$\Psi_{\text{DM}}^1(X) = \Psi_{\text{rw}}(X), \quad \Lambda_{\text{DM}} = I - \Lambda_{\text{rw}}.$$

Hence, the largest eigenvalue is $\lambda_0 = 1$ and the corresponding eigenvector is $\mathbf{1} \in \mathbb{R}^{m_*}$, i.e., $A_{\text{rw}}\mathbf{1} = \mathbf{1}$. Similar to the Laplacian eigenmaps, this constant vector is useless for the embedding. Hence, we solve (11.14) for the $(s+1)$ -dimensional embedding $\Psi(X)$, and use its 2nd to the last rows and the eigenvalues $\lambda_1 \geq \dots \geq \lambda_s$ for $\Psi_{\text{DM}}(X)$ in (11.15).

Two important remarks are in order here. First, one can also define the symmetric version $\tilde{A}_{\text{sym}} := \tilde{D}^{-\frac{1}{2}} \tilde{W} \tilde{D}^{-\frac{1}{2}}$ in a similar manner as L_{sym} of (11.11). In order to compute $\Psi_{\text{DM}}(X)$, we can use either: 1) SVD of \tilde{A}_{rw} and use the 2nd to $(s+1)$ st left singular vectors; or 2) symmetric eigenvalue solver for \tilde{A}_{sym} , and then normalize its 2nd to $(s+1)$ st eigenvectors by the 1st eigenvector $\tilde{D}^{\frac{1}{2}}\mathbf{1}$ corresponding to the eigenvalue $\lambda_0 = 1$ (i.e., $\tilde{A}_{\text{sym}}\tilde{D}^{\frac{1}{2}}\mathbf{1} = \tilde{D}^{\frac{1}{2}}\mathbf{1}$). For the details on this normalization, see Appendix B of [19].

Second, Coifman and Lafon defined the approximate *diffusion distance* between two data points \mathbf{x}_i and \mathbf{x}_j in the original ambient space, which is simply the ℓ^2 distance between the embedded points in the diffusion space, i.e.,

$$\delta_{\text{DM}}^t(\mathbf{x}_i, \mathbf{x}_j) := \|\Psi_{\text{DM}}^t(\mathbf{x}_i) - \Psi_{\text{DM}}^t(\mathbf{x}_j)\| = \sqrt{\sum_{k=1}^s \lambda_k^{2t} (\psi_k(\mathbf{x}_i) - \psi_k(\mathbf{x}_j))^2},$$

where $\psi_k(\mathbf{x}_i)$ is the (k, i) th entry of the matrix $\Psi_{\text{DM}}(X) \in \mathbb{R}^{s \times m_*}$ following the convention of (11.1), which does not include the first eigenvector $\mathbf{1}$.

11.2.4 Out-Of-Sample Multiscale Extension Via Geometric Harmonics

Our goal is in classifying newly obtained unlabeled test ensembles based on a classification rule learned from the labeled training ensembles. In order to make meaningful inference from the training ensembles to the test ensembles, we need to have the same low-dimensional representation system for all ensembles. That is, we need to embed test ensembles into the same reduced space as the training ensembles. Hence, it becomes necessary for us to extend the embedding map computed

on the training ensembles to the test ensembles. If the embedding map is PCA/CS, then such extension is quite simple: it suffices to expand the centered test data matrix $\tilde{Y} = YH$ by the left singular vectors U_s of \tilde{X} , i.e., $U_s^T \tilde{Y}$ is the extension of the test ensembles by PCA. For a nonlinear map such as the Laplacian eigenmaps and the diffusion maps, we employ the multiscale extension scheme proposed in [19], which is based on “geometric harmonics” originally introduced in [21, Chap. 3] and [10]. Let us call this scheme GHME (geometric harmonics multiscale extension) for short. We now review the GHME scheme using our matrix-vector notation as in the previous sections.

The GHME scheme is an improvement of the *Nyström extension method* proposed in [5, 14]. First, consider the Gaussian kernel matrix defined on the training dataset \mathcal{X} with the scale parameter $\sigma > 0$, which is different from ε used in the weight function in (11.7) for constructing the Laplacian eigenmaps and the diffusion maps as follows:

$$W_\sigma(X) := (w_\sigma(\mathbf{x}_i, \mathbf{x}_j)) = \left(e^{-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / \sigma^2} \right) \in \mathbb{R}^{m_* \times m_*}. \quad (11.16)$$

Since $W_\sigma(X)$ is positive semi-definite, let us consider the eigenfunction expansion of this kernel matrix:

$$W_\sigma(X) = \Phi^T M \Phi, \quad \Phi := \begin{bmatrix} \phi_1^T \\ \vdots \\ \phi_{m_*}^T \end{bmatrix} \in \mathbb{R}^{m_* \times m_*}, \quad M := \text{diag}(\mu_1, \dots, \mu_{m_*}). \quad (11.17)$$

where the nonnegative eigenvalues $\{\mu_j\}$ are sorted in non-increasing order, and the column vectors of Φ^T , i.e., $\{\phi_1, \dots, \phi_{m_*}\}$ form an orthonormal basis for $\ell^2(\mathcal{X}) \subset \mathbb{R}^{m_*}$, and $\phi_i := (\phi_i(\mathbf{x}_1), \dots, \phi_i(\mathbf{x}_{m_*}))^T$. Now, consider the k th eigenpair (μ_k, ϕ_k) , i.e., $W_\sigma(X)\phi_k = \mu_k\phi_k$. The i th row of this equality gives us

$$\phi_k(\mathbf{x}_i) = \frac{1}{\mu_k} \sum_{j=1}^{m_*} w_\sigma(\mathbf{x}_i, \mathbf{x}_j) \phi_k(\mathbf{x}_j).$$

The Nyström extension of ϕ_k from \mathcal{X} to $\mathbf{y} \in \mathcal{Y}$ is defined as

$$\bar{\phi}_k(\mathbf{y}) := \frac{1}{\mu_k} \sum_{j=1}^{m_*} w_\sigma(\mathbf{y}, \mathbf{x}_j) \phi_k(\mathbf{x}_j). \quad (11.18)$$

Since the eigenvectors $\{\phi_k\}$ form an orthonormal basis for $\ell^2(\mathcal{X})$, any function $\mathbf{f} := (f(\mathbf{x}_1), \dots, f(\mathbf{x}_{m_*}))^T \in \ell^2(\mathcal{X})$ (e.g., \mathbf{f}^T can be any row of the Laplacian eigenmap $\Psi_{\text{rw}}(X)$ or the diffusion map $\Psi_{\text{DM}}^l(X)$) can be expanded as

$$\mathbf{f} = \sum_{k=1}^{m_*} \langle \mathbf{f}, \phi_k \rangle \phi_k,$$

and the j th entry of both sides gives us

$$f(\mathbf{x}_j) = \sum_{k=1}^{m_*} \langle \mathbf{f}, \boldsymbol{\phi}_k \rangle \phi_k(\mathbf{x}_j).$$

Thus, the Nyström extension of f from \mathcal{X} to $\mathbf{y} \in \mathcal{Y}$ can be defined as

$$\bar{f}(\mathbf{y}) := \sum_{k=1}^{m_*} \langle \mathbf{f}, \boldsymbol{\phi}_k \rangle \bar{\phi}_k(\mathbf{y}).$$

In order to understand what we have done here, let us plug (11.18) in the righthand side of the above equation.

$$\begin{aligned} \bar{f}(\mathbf{y}) &= \sum_{k=1}^{m_*} \frac{\langle \mathbf{f}, \boldsymbol{\phi}_k \rangle}{\mu_k} \sum_{j=1}^{m_*} w_{\sigma}(\mathbf{y}, \mathbf{x}_j) \phi_k(\mathbf{x}_j) \\ &= \sum_{k=1}^{m_*} \frac{\boldsymbol{\phi}_k^T \mathbf{f}}{\mu_k} w_{\sigma}(\mathbf{y}, :) \boldsymbol{\phi}_k \\ &= w_{\sigma}(\mathbf{y}, :) \boldsymbol{\Phi}^T M^{-1} \boldsymbol{\Phi} \mathbf{f}, \end{aligned} \tag{11.19}$$

where we borrowed the convenient notation of MATLAB[®] 2, i.e.,

$$w_{\sigma}(\mathbf{y}, :) := (w_{\sigma}(\mathbf{y}, \mathbf{x}_1), \dots, w_{\sigma}(\mathbf{y}, \mathbf{x}_{m_*})) \in \mathbb{R}^{1 \times m_*}.$$

We observe that the range of the extension in (11.18) is proportional to σ . If the ratio $\|\mathbf{y} - \mathbf{x}_j\|/\sigma$ is large for all $\mathbf{x}_j \in \mathcal{X}$, then $\bar{\phi}_k(\mathbf{y})$ will be numerically small and hence may not be meaningful. Hence, the extension scale σ should be as large as possible. On the other hand, for large enough σ , the kernel matrix $W_{\sigma}(X)$ defined in (11.16) becomes ill-conditioned, i.e., μ_k tends to 0 more quickly compared to the case of small σ . Thus, the Nyström extension in (11.18) will blow up. Furthermore, it is well known that the extension range depends on the smoothness of the function to be extended [21, Chap. 3], [10]. If f is fairly smooth, it can be extended far away from the training set while it has limited extension range if f varies wildly on \mathcal{X} . We can clearly see these problems in (11.19): it truly boils down to the ill-conditionedness of $W_{\sigma}(X)$. In fact, if we can compute M^{-1} without blowing up, i.e., $\mu_{m_*} \gtrless 0$, then $\boldsymbol{\Phi}^T M^{-1} \boldsymbol{\Phi} = W_{\sigma}(X)^{-1}$ because of (11.17). Moreover, by setting $\mathbf{y} = \mathbf{x}_j \in \mathcal{X}$ in the righthand side of (11.19) recovers $f(\mathbf{x}_j)$:

$$w_{\sigma}(\mathbf{x}_j, :) \boldsymbol{\Phi}^T M^{-1} \boldsymbol{\Phi} \mathbf{f} = w_{\sigma}(\mathbf{x}_j, :) W_{\sigma}(X)^{-1} \mathbf{f} = \mathbf{e}_j^T \mathbf{f} = f(\mathbf{x}_j),$$

where $\mathbf{e}_j \in \mathbb{R}^{m_*}$ is the j th standard basis vector in \mathbb{R}^{m_*} . In practice, however, $W_{\sigma}(X)$ is ill-conditioned, and we need to truncate M^{-1} to the first $p \times p$ submatrix and $\boldsymbol{\Phi}$

² MATLAB is a registered trademark of The MathWorks, Inc., 3 Apple Hill Drive, Natick, MA 01760-2098.

to the first p rows where p satisfying $1 \leq p < m_*$ must be appropriately chosen. Let M_p^{-1} , Φ_p be these truncated matrices. Then, (11.19) can be approximated without blowup:

$$\begin{aligned}\bar{f}(\mathbf{y}) &\approx w_\sigma(\mathbf{y}, :) \Phi_p^T M_p^{-1} \Phi_p \mathbf{f} \\ &= w_\sigma(\mathbf{y}, :) W_{\sigma,p}(X)^\dagger \mathbf{f},\end{aligned}\tag{11.20}$$

where $W_{\sigma,p}(X)^\dagger := \Phi_p^T M_p^{-1} \Phi_p$ is the *pseudoinverse* of $W_\sigma(X)$ using their top p singular values and vectors. Hence, if we want to extend a low-dimensional embedding map $\Psi(X) = [\boldsymbol{\psi}_1 | \cdots | \boldsymbol{\psi}_s]^T$, we have

$$\begin{aligned}\bar{\Psi}(\mathbf{y}) &= \begin{bmatrix} \bar{\boldsymbol{\psi}}_1(\mathbf{y}) \\ \vdots \\ \bar{\boldsymbol{\psi}}_s(\mathbf{y}) \end{bmatrix} \approx \{w_\sigma(\mathbf{y}, :) \Phi_p^T M_p^{-1} \Phi_p [\boldsymbol{\psi}_1 | \cdots | \boldsymbol{\psi}_s]\}^T \\ &= \{w_\sigma(\mathbf{y}, :) \Phi_p^T M_p^{-1} \Phi_p \Psi(X)^T\}^T \\ &= \Psi(X) \Phi_p^T M_p^{-1} \Phi_p w_\sigma(:, \mathbf{y}) \\ &= \Psi(X) W_{\sigma,p}(X)^\dagger w_\sigma(:, \mathbf{y}),\end{aligned}$$

where $w_\sigma(:, \mathbf{y}) := (w_\sigma(\mathbf{x}_1, \mathbf{y}), \dots, w_\sigma(\mathbf{x}_{m_*}, \mathbf{y}))^T = w_\sigma(\mathbf{y}, :)^T \in \mathbb{R}^{m_* \times 1}$.

The GHME scheme determines this rank p of the pseudoinverse by the following:

$$p = \arg \max_{1 \leq k \leq m_*} \left\{ \frac{\mu_1}{\mu_k} \leq \eta \right\}.\tag{11.21}$$

where $\eta > 0$ is some fixed condition number. In other words, p is the largest possible stable rank of $W_\sigma(X)$, which is bounded from above by η . Clearly, if one sets η too large, $p = m_*$ may occur. The extension \bar{f} in (11.20) is well-defined on $\mathcal{X} \cup \mathcal{Y}$, but it is not equal to f on the training set \mathcal{X} unless $p = m_*$ and σ is set so small that $W_\sigma(X)$ has a stable inverse. Such a case, however, is not of our interest because setting σ too small practically disconnects data points in \mathcal{X} . In fact as $\sigma \rightarrow 0$, $W_\sigma(X) \rightarrow I$ as long as $\mathbf{x}_i \neq \mathbf{x}_j$ for all $i \neq j$ in \mathcal{X} . Yet observe that if the value of σ decreases, the eigenvalues $\mu_k \rightarrow 0$ more slowly. This allows us to use larger p in (11.20), making \bar{f} a better approximation of f on \mathcal{X} . Based on this observation, the GHME iteratively searches for an extension \bar{f} that approximates f on \mathcal{X} with an preset error tolerance $\rho > 0$ by slowly decreasing the value of the extension scale σ . The GHME scheme can be summarized as follows:

Algorithm 11.2.1 (The GHME of Lafon, Keller, and Coifman[19]). Suppose f is a function defined on the training set \mathcal{X} and to be extended to a new dataset \mathcal{Y} .

Step 1: Fix a condition number $\eta > 0$ and an error tolerance $\rho > 0$. Set the extension scale $\sigma = \sigma_0$ for some large value σ_0 .

Step 2: Compute eigenvalues $\{\mu_k\}$ and orthonormal eigenvectors $\{\phi_k\}$ of the Gaussian kernel matrix $W_\sigma(X)$ and expand f (on the training set \mathcal{X}) in this eigenbasis.

Step 3: On the training set \mathcal{X} , approximate f by \bar{f} defined in (11.20) by finding p in (11.21). Then compute the approximation error

$$Err := \left(\sum_{k>p} |\langle f, \phi_k \rangle|^2 \right)^{1/2}.$$

If $Err > \rho$, set $\sigma \leftarrow \frac{1}{2}\sigma$ and return to Step 2. Otherwise, continue.

Step 4: Using the value of p obtained in Step 3, compute the final approximate extension defined in (11.20) for each $\mathbf{y} \in \mathcal{Y}$:

$$\bar{f}(\mathbf{y}) \approx w_\sigma(\mathbf{y}, :) \Phi_p^T M_p^{-1} \Phi_p \mathbf{f} = w_\sigma(\mathbf{y}, :) W_{\sigma,p}(X)^\dagger \mathbf{f}.$$

11.3 Earth Mover's Distance

Once our training ensembles are embedded in a lower-dimensional space, say \mathbb{R}^s (normally $1 \leq s \ll \min(d, m_*)$), via the diffusion map or Laplacian eigenmaps so as the test ensembles are via the GHME scheme as reviewed in the previous section, we are ready to classify the latter. To do so, we need to quantify a discrepancy (or measure a “distance”) between a test ensemble and each training ensemble in \mathbb{R}^s . One of the simplest ideas such as the sum of the pairwise Euclidean distances (in the embedding space) between the data points in one ensemble and those in the other ensemble will not work because: (1) there may be a few outliers in those ensembles that ruin such a distance measure and (2) the number of data points (signals) in each ensemble may be different in general so that the simple sum of the distances may lead to an erroneous label assignment.

Another idea is to use the Hausdorff distance (HD), which was used by Lafon, Keller, and Coifman as the ensemble distance measure [19]. The HD between any two ensembles $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^s$ is defined as

$$d_H(\mathcal{X}, \mathcal{Y}) := \max \left(\max_{\mathbf{y} \in \mathcal{Y}} \min_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|, \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \mathcal{Y}} \|\mathbf{x} - \mathbf{y}\| \right), \quad (11.22)$$

where $\|\cdot\|$ denotes the Euclidean distance in \mathbb{R}^s . As our numerical experiments in Sect. 11.5 demonstrate, also as one can easily imagine from this definition, the HD is still quite sensitive to the outliers.

The above considerations have led us to adopt the *Earth Mover's Distance* (EMD) as a discrepancy/distance measure between ensembles, which is more robust and more suitable in our classification problems than the other measures. In this section, we briefly review the key aspects of the EMD.

The definition of EMD is based on the solution to a discrete *optimal mass transportation problem*. EMD represents the minimum cost of moving earth (or sand) from some source locations to fill up holes at some sink locations. In other words, given any two mass (or probability) distributions, one of them can be viewed as a distribution of earth and the other a distribution of holes, then the EMD between the two distributions is the minimum cost of rearranging the mass in one distribution to obtain the other. In the continuous setting, this problem is known as the *Monge–Kantorovich optimal mass transfer problem* and has been well studied over the past 100 years (see [13] for an introductory reading on the problem). The importance here is that EMD can be used to measure the discrepancy between two multidimensional distributions.

In the discrete setting, the optimal mass transfer problem can be formulated as a linear optimization problem as follows [33, 34]. Suppose we want to quantify the discrepancy between two ensembles of signals (or feature vectors) of length $s \in \mathbb{N}$. Suppose all the signals in one ensemble are grouped into m clusters using a clustering algorithm such as the K -means algorithm (see Sect. 11.4.1 for what we actually use to construct a signature in our experiments; for the details of various clustering algorithms, see e.g., [17, Chap. 14]) and this ensemble is represented by a set of m cluster centers and the associated weights (member population within each cluster) as $P = \{(\hat{\mathbf{x}}_1, p_1), \dots, (\hat{\mathbf{x}}_m, p_m)\}$. Note that one can optionally normalize the weights p_i 's so that $\sum_i p_i = 1$, which allows us to view P as a probability distribution of the signals in \mathcal{X} over \mathbb{R}^s . Suppose the other ensemble \mathcal{Y} is represented, using the same procedure, by a set of n cluster centers and weights $Q = \{(\hat{\mathbf{y}}_1, q_1), \dots, (\hat{\mathbf{y}}_n, q_n)\}$. Such P and Q are called the *signatures* of the ensembles. Viewing P as a source (earth) distribution and Q as a sink (hole) distribution, we can now formulate the discrete optimal mass transport problem leading to the EMD as a measure of the discrepancy between two signatures P and Q .

Suppose the cost of moving one unit of mass from $\hat{\mathbf{x}}_i$ to $\hat{\mathbf{y}}_j$ is $c(\hat{\mathbf{x}}_i, \hat{\mathbf{y}}_j)$, and f_{ij} denotes the amount of mass flow from $\hat{\mathbf{x}}_i$ to $\hat{\mathbf{y}}_j$. There are many possibilities in defining this cost, but we use $(1/2)\|\hat{\mathbf{x}}_i - \hat{\mathbf{y}}_j\|^2$, i.e., the half of the *squared* Euclidean distance, which gives more preference to closely placed points than the usual Euclidean distance between them. Then, the transportation cost is defined as

$$\text{COST}(P, Q, \mathbf{F}) := \sum_{i=1}^m \sum_{j=1}^n c(\hat{\mathbf{x}}_i, \hat{\mathbf{y}}_j) f_{ij},$$

where $\mathbf{F} := [f_{ij}] \in \mathbb{R}^{m \times n}$. The optimal mass transfer problem seeks the flow \mathbf{F}^* that transfers the maximum allowable amount of earth to fill up the holes with minimum total transportation cost, i.e.,

$$\mathbf{F}^* = \arg \min_{\mathbf{F} \in S} \text{COST}(P, Q, \mathbf{F}),$$

where $\mathbf{F} \in S \subset \mathbb{R}^{m \times n}$ means that \mathbf{F} must satisfy the following set of constraints:

- (a) $f_{ij} \geq 0$, for all i, j .
- (b) $\sum_{j=1}^n f_{ij} \leq p_i$, for all $1 \leq i \leq m$.

- (c) $\sum_{i=1}^m f_{ij} \leq q_j$, for all $1 \leq j \leq n$.
 (d) $\sum_{i=1}^m \sum_{j=1}^n f_{ij} = \min \left(\sum_{i=1}^m p_i, \sum_{j=1}^n q_j \right)$.

These constraints ensure that: (a) one can only move earth from P to Q , not vice versa; (b) the amount of earth moved from P is no more than the sum of the weights p_i ; (c) the amount of earth received at Q is no more than the sum of the weights q_j ; and (d) the maximum allowable amount of earth is moved.

Once the optimal flow \mathbf{F}^* from P to Q is found, EMD is then defined as the total cost normalized by the total flow:

$$\text{EMD}(P, Q) := \frac{\text{COST}(P, Q, \mathbf{F}^*)}{\sum_{i=1}^m \sum_{j=1}^n f_{ij}^*} = \frac{\sum_{i=1}^m \sum_{j=1}^n c(\hat{\mathbf{x}}_i, \hat{\mathbf{y}}_j) f_{ij}^*}{\sum_{i=1}^m \sum_{j=1}^n f_{ij}^*}. \quad (11.23)$$

Notice that the normalization factor is the total weight of the smaller signature due to the constraint (d). This normalization ensures that smaller signatures are not favored in the case when two signatures have different total weights. Furthermore, EMD is symmetric, i.e., $\text{EMD}(P, Q) = \text{EMD}(Q, P)$ for any two distributions P and Q . Notice also that each signature may consist of different number of clusters or “bins” in the reduced embedding space as indicated by m and n above. Since there is no guarantee to have the same number of cluster centers in the signature for each ensemble, EMD is more suitable for our classification problem compared to the other metrics. Finally, notice that the EMD directly depends on P and Q , which are not uniquely determined from the ensembles \mathcal{X} and \mathcal{Y} . In other words, the EMD depends on the clustering algorithm to construct P and Q . Hence, it requires care to construct the signatures, which we will discuss in detail in Sect. 11.4.1.

11.4 An Algorithm for Signal Ensemble Classification Using Low-Dimensional Embeddings with Earth Mover’s Distance

In this section, we will describe our proposed algorithm to classify test signal ensembles $\mathcal{Y}^1, \dots, \mathcal{Y}^N$ given training signal ensembles $\mathcal{X}^1, \dots, \mathcal{X}^M$ using a low-dimensional embedding and the EMD as the ensemble distance measure. Then, we will describe how to construct appropriate signatures of ensembles in the embedding space. Finally, we will discuss an important practical issue, i.e., how to select several key tuning parameters in our algorithm.

We summarize our proposed method for signal ensemble classification in the following algorithm.

Algorithm 11.4.1 (Signal Ensemble Classification via a Low-Dimensional Embedding and EMD). Let $\mathcal{X} = \bigcup_{i=1}^M \mathcal{X}^i$, $\mathcal{Y} = \bigcup_{j=1}^N \mathcal{Y}^j$ be the training and test ensembles, respectively.

Step 1. Training Stage using \mathcal{X} :

- i. Preset a large enough initial dimension $1 < s_0 < \min(d, m_*)$ of the embedding space.
- ii. Choose Ψ from $\{\Psi_{\text{PCA}}, \Psi_{\text{rw}}, \Psi_{\text{sym}}, \Psi_{\text{DM}}^1\}$, construct a low-dimensional embedding map $\Psi: \mathcal{X} \subset \mathbb{R}^d \rightarrow \mathbb{R}^{s_0}$, and embed \mathcal{X} into the temporary reduced space in \mathbb{R}^{s_0} .
- iii. For $i = 1 : M$, construct a signature P^i using $\Psi(X^i) \in \mathbb{R}^{s_0}$.
- iv. Determine the appropriate reduced dimension $1 \leq s \leq s_0$.
- v. For $i = 1 : M$, adjust the signature P^i in the final reduced space \mathbb{R}^s .

Step 2. Test Stage for \mathcal{Y} :

- i. Extend the learned map Ψ in Step 1 to the test ensembles \mathcal{Y} .
- ii. For $j = 1 : N$, construct the signature Q^j of \mathcal{Y}^j .
- iii. For $j = 1 : N$, compute $\{\text{EMD}(P^i, Q^j)\}_{i=1}^M$ and find $i_j := \arg \min_{1 \leq i \leq M} \text{EMD}(P^i, Q^j)$.
Assign the test ensemble \mathcal{Y}^j the label of \mathcal{X}^{i_j} . That is, **assign a label by the 1-nearest neighbor classifier with EMD as its base distance measure.**

We will now describe what are important for implementing and running the above algorithm for given signal ensembles.

11.4.1 Signature Construction and Embedding Dimension Estimation

The above overall algorithm did not show the detailed differences between the case when one chooses Ψ_{PCA} and the cases of the nonlinear embedding maps, i.e., Ψ_{rw} , Ψ_{sym} , and Ψ_{DM}^1 . In this section, we will describe these differences in detail.

Let us first describe Step 1-iii, i.e., the signature construction step. If we use PCA as Ψ , then we use the standard K -means algorithm to construct the clusters in the top s_0 PCA coordinates of \mathcal{X}^i . For the K -means algorithm, it is essential to determine the number of clusters K_i for each training ensemble \mathcal{X}^i . There are many possibilities for determining the number of optimal clusters (see e.g., [17, Sect. 14.3.11]), out of which we use a heuristic yet simple method called the “elbow” criterion. This can be explained as follows. As we increase the number of clusters, the fitness (or within-cluster dissimilarity) function, which is the sum of all point-to-centroid distances, will decrease very rapidly at first then slowly. This will result in a “kink” or “elbow” in the plot of the fitness function versus the number of clusters. Hence, we choose the number of clusters at the elbow as the natural candidate for the number of clusters to form for the given data.

On the other hand, if we use the Laplacian eigenmaps or the diffusion map as Ψ , then we use the *elongated* K -means (*ekmeans*) algorithm [39] to determine the number of clusters in the signature for each training ensemble \mathcal{X}^i . The idea of *ekmeans* was adapted from the spectral clustering algorithm proposed in [30] by

replacing the Euclidean distance with an elongated distance in the computation of point-to-center distances. In [39], the elongated distance (e-dist) between two points $\mathbf{x}, \mathbf{c} \in \mathbb{R}^{s_0}$ is defined as

$$\text{e-dist}(\mathbf{x}, \mathbf{c}) := (\mathbf{x} - \mathbf{c})^T M_\alpha (\mathbf{x} - \mathbf{c}), \quad M_\alpha := \frac{1}{\alpha} (I - P_{\mathbf{c}}) + \alpha P_{\mathbf{c}}, \quad P_{\mathbf{c}} := \frac{\mathbf{c}\mathbf{c}^T}{\mathbf{c}^T \mathbf{c}}. \quad (11.24)$$

As one can see from this definition, the e-dist considers two components of the vector $\mathbf{x} - \mathbf{c}$: (1) the radial component $P_{\mathbf{c}}(\mathbf{x} - \mathbf{c})$, i.e., the orthogonal projection of $\mathbf{x} - \mathbf{c}$ onto the direction \mathbf{c} ; and (2) the traversal component $(I - P_{\mathbf{c}})(\mathbf{x} - \mathbf{c})$, i.e., the complementary projection of $\mathbf{x} - \mathbf{c}$ with respect to \mathbf{c} . The parameter α controls the balance between these two components: the smaller α becomes, the more traversal direction is emphasized in the distance computation. In all of our numerical experiments below, we set $\alpha = 0.2$, the value recommended by the authors in [39]. The use of the e-dist allows us to group points lying inside a thin elongated ellipsoid to form a cluster, as opposed to inside a sphere.

Ekmeans exploits the geometric properties of the eigenvectors of the graph Laplacian matrix or the adjacency matrix to cluster the data and automatically determine the number of intrinsic clusters. It starts the clustering process in the top 2 eigenspace with three initialized centers: two centers at two different elongated clusters and one at the origin of the eigenspace. If there are more than two elongated clusters, the center at the origin will be dragged to a cluster not yet accounted for. Then the algorithm moves the clustering process to the top 3 eigenspace, creating three cluster centers and adds a fourth center at the origin. This process is repeated until no additional cluster is found. This clustering process stops at the top K eigenspace if there are K (intrinsic) clusters in the data.

Step 1-iv determines the appropriate dimension s of the final reduced embedding space. In the case of $\Psi = \Psi_{\text{PCA}}$, s is determined by the decay of the singular values σ_k of the centered training data matrix \tilde{X} . In fact, using PCA, we can combine Steps 1-iii, iv, and v into a single step by first estimating the effective dimension s by the decay of the singular values, then construct a signature in \mathbb{R}^s for each \mathcal{X}^i using the K -means algorithm with the elbow criterion. By this way, we do not need to readjust the signatures in Step 1-v. As for the Laplacian eigenmaps and diffusion maps, suppose the *ekmeans* algorithm finds an intrinsic dimension K_i for each training ensemble \mathcal{X}^i . This number K_i is also the intrinsic number of clusters in that ensemble. This intrinsic number of K_i clusters does not change when the ensemble \mathcal{X}^i is embedded into an eigenspace of dimension greater than K_i . This is because no additional cluster other than the K_i elongated clusters already accounted for will be found even if \mathcal{X}^i is embedded into such a larger eigenspace. Therefore, it is natural to set the dimension of the reduced embedding space for the training ensembles \mathcal{X} to be the maximum of K_i over $i = 1, \dots, M$, i.e., $s := \max_{1 \leq i \leq M} K_i$.

Step 1-v is necessary for both the Laplacian eigenmaps and the diffusion map as Ψ because of the following reason. The cluster centers for ensemble \mathcal{X}^i determined by *ekmeans* in Step 1-iii are vectors in the top K_i -dimensional subspace of the temporary reduced space \mathbb{R}^{s_0} . Step 1-iv determines the final dimension s of the

reduced space. In order for us to use EMD in Step 2, however, all the cluster centers must be in \mathbb{R}^s . Hence, to bring all cluster centers into \mathbb{R}^s , we need to re-cluster \mathcal{X}^i 's embedded points in \mathbb{R}^s by running the standard K -means with the e-dist as its base distance measure to reform the K_i clusters. Here, we use the previous cluster memberships as a starting condition for the standard K -means. At the end of this re-clustering process, all cluster centers of the signature for each training ensemble are in the final reduced space \mathbb{R}^s .

Finally, in Step 2-i, we extend the learned map to the test ensembles. If $\Psi = \Psi_{\text{PCA}}$, then this extension can be computed simply by $U_s^T \tilde{Y}$. If $\Psi = \Psi_{\text{rw}}, \Psi_{\text{sym}}$, or Ψ_{DM} , then we use the GHME Algorithm 11.2.1.


11.4.2 On Tuning Parameters

Other than the dimensionality s of the reduced embedding space, there are three more tuning parameters to be determined in our proposed method: the scale ε for the weights $w_\varepsilon(\mathbf{x}_i, \mathbf{x}_j)$ defined in (11.7); the error tolerance ρ in approximating the extension of the nonlinear embedding maps; and the cutoff bound η for the condition number of the extension kernel. Clearly, we have to select each of these parameters wisely. We describe now how to tune these parameters.

11.4.2.1 Determination of ε

The scale $\varepsilon > 0$ for Gaussian weights $w_\varepsilon(\mathbf{x}, \mathbf{y})$ should be chosen so that the corresponding weighted graph is numerically connected. Connectedness of the graph guarantees the existence and uniqueness of the eigenvector corresponding to the smallest eigenvalue 0 of L_{rw} in (11.10) and the largest eigenvalue 1 of \tilde{A}_{rw} in (11.14), and therefore of the stationary distribution of the Markov random walk process on the graph. Thus, when computing diffusion maps, the value for ε must be large enough to ensure that every point in the graph is connected to at least one other point. It is clear, however, that any affinity or dissimilarity between the data points is obscured if ε is too large since w_ε converges to 1 regardless of its arguments (i.e., every data point is equally similar to any other point) as ε increases to infinity. One suggestion given in [23] is to choose ε “in the order of the mean distance of a point to its k th nearest neighbor” (where k is proportional to the logarithm of the number of points in the graph). In our numerical experiments below, we select ε to be the mean of the Euclidean distances from each point to its k -nearest neighbors. We determine this k by running cross-validation experiments on the training set \mathcal{X} . We start with $k = 1$, run a cross-validation trial, and record the results. Next, continue to increment k by Δk and run a cross-validation trial until the classification results stop to improve. When this happens, the previous value for k , say k^* , may already be optimal. To ensure that this k^* is indeed optimal, however, we decrease Δk to $\lfloor \Delta k/2 \rfloor$ and search for possibly better value for k around k^* if we set Δk large. The initial

Fig. 11.1 The largest 300 eigenvalues of the diffusion kernel in one trial of the lip-reading experiments ($\varepsilon = 657$)



size of Δk depends on the total number of the training signals m_* . In our lip-reading example, we start with $\Delta k = 5$ since the training set is large ($990 \leq m_* \leq 1,204$). However, in the classification of underwater objects, m_* is small ($96 \leq m_* \leq 120$), so we start with $\Delta k = 1$. As an illustration, in our lip-reading example below, the average value for k^* determined by this cross-validation search method (over 100 experimental trials) turned out to be 16 when EMD was used for the ensemble distance measure and 14 with HD. These are approximately 1.4%–1.6% of the training signals m_* . For the underwater object classification, any values of $k^* \in [2, 30]$ with EMD and $k^* \in [8, 40]$ with HD produced the same optimal classification results. These k^* 's correspond to about 2%–40% of m_* . We note that when k is, say, 5% of the number of the training signals, the corresponding Gaussian-weighted graph has approximately 5% of its edges weighted by a weight greater than or equal to $1/e$, and the rest have smaller weights. In other words, the graph is sparse but not too sparse. With this choice of ε , the spectrum of the diffusion kernel decays relatively fast. In Fig. 11.1, we plot the largest 300 eigenvalues of the normalized diffusion matrix \tilde{A}_{rw} defined in (11.13) of one trial of the lip-reading experiments. The value of ε is 657 corresponding to $k^* = 16$. We see that the eigenvalues decrease quickly. Fast decay of the spectrum implies that any random walk initiated on the graph converges quickly to steady state and that the diffusion distance can be approximated more accurately with a smaller number of eigenfunctions. In other words, we should be able to detect clustering behaviors in the data with a small number of time steps t for such a case. In the numerical experiments in Sect. 11.5, we used $t = 1$ in the diffusion maps defined in (11.15). Setting $t = 1$ and optimizing ε is certainly an easier strategy compared to the simultaneously seeking the optimal combination of $t > 1$ and ε although the latter may generate better classification results. Also, there may be datasets on which some value of $t > 1$ results in better classification than setting $t = 1$. At least for those datasets in Sect. 11.5, however, having found the value ε appropriate for the data while setting $t = 1$ was enough for identifying grouping patterns.

We note that the value of ε or k found to be optimal for the diffusion maps may not be optimal for the Laplacian eigenmaps, and vice versa. Thus, we need to determine the optimal k separately for each case. We could start at $k = 1$ and increment by Δk as described above. However, if we have already found the optimal value k^* for the diffusion map (or the Laplacian eigenmap) and if we want to save the computational time to find the optimal value k for the Laplacian eigenmap (or the diffusion map), then we could proceed as follows. We start the search with $k = k^*$; then we increase k by e.g., $\Delta k = 5$ and proceed as described above. If the classification is not improving, we also need to search in the opposite direction, i.e., decrease k by $\Delta k = 5$ and proceed (in decreasing direction) as described above. In our numerical experiments below, we computed the diffusion map first. Therefore, when searching for the optimal k for the Laplacian eigenmap, we took the approach described in this paragraph. The optimal value k found when computing the Laplacian eigenmaps (the random walk version) is 60 (EMD) and 67 (HD) for the lip-reading experiments (averaging over 100 experimental trials) while in the underwater object classification any values of $k \in [12, 40]$ (EMD) and $k = 7$ or 13 (HD) generated the optimal classification results. For the symmetric version of the Laplacian eigenmaps, the average optimal k was 57 (EMD) and 62 (HD) for the lip-reading experiments while any values of $k \in [3, 8]$ (EMD) and $k \in [1, 8]$ (HD) were optimal for the underwater object classification.

11.4.2.2 Determination of ρ

When choosing a value for the error tolerance ρ for the approximation of the out-of-sample extension of the nonlinear embedding maps in Algorithm 11.2.1 that is used in Step 2-i of our Algorithm 11.4.1, we should keep in mind that a small error bound ρ corresponds directly to a small extension scale σ as we discussed in Sect. 11.2.4. Suppose we know a priori that the function we want to extend is fairly smooth on our training ensembles \mathcal{X} . Then we can expect the extension to have a large extension range. In this case, we can be greedy and set ρ small, i.e., we want the extension $\bar{\Psi}$ to be very close to Ψ on \mathcal{X} . A heuristic value to set for ρ is 1% of the number of the training signals m_* . This gives an average of 0.01 bound on the error at each point where the approximation is being computed.

11.4.2.3 Determination of η

To determine a cutoff threshold η for the condition number of the Gaussian extension kernel matrix $W_\sigma(X)$ in (11.16), we have to keep in mind the approximation error tolerance ρ . If ρ is small, then η has to be large. In addition, as σ increases, the condition number of the kernel matrix also increases. To predict how large η might get, we can take advantage of Step 1 of our Algorithm 11.4.1. Let κ be the condition number of $\tilde{A}_{\text{rw}}, \tilde{A}_{\text{sym}}, L_{\text{rw}},$ or L_{sym} with the parameter ε optimally chosen in Step 1. It is easy to show that the condition number of $W_\sigma(X)$ is proportional to κ

if $\sigma = \varepsilon$. Furthermore, since we usually set the initial $\sigma_0 > \varepsilon$ in Algorithm 11.2.1, the condition number of the initial Gaussian kernel is larger than κ . Hence, we can consider setting η larger than κ and inversely proportional to ρ . In our numerical experiments, we set $\eta = \min(\kappa, 10^5)/\rho$.

11.5 Numerical Experiments and Results

We now illustrate how our proposed algorithm can be applied to signal ensemble classification problems where the data characterizing each object consist of ensembles of signals instead of a single signal. We will show two examples of application. The first example is classification of underwater objects via analyzing Synthetic Aperture Sonar (SAS) waveforms reflected from the objects. The second example is a lip-reading application in which we identify the spoken word from a sequence of video frames extracted from a silent video segment. As we mentioned earlier, we use PCA, Laplacian eigenmaps (two different normalizations L_{TW} and L_{sym}), and diffusion map. We will also compare the performance of EMD with that of the Hausdorff distance (HD) defined in (11.22) in our classification problems.

11.5.1 Classification of Underwater Objects

The data in this example were provided by the Naval Surface Warfare Center, Panama City (NSWC-PC), FL. They were collected from three different controlled experiments in a fresh water test pond at NSWC-PC. For details of the experiments, see [29]. In each of the three experiments, two objects were placed – either buried in the sand or proud – at the bottom of the pond. In each experiment, one object was a sphere made of an iron casing filled with a different material, and the other object was a solid aluminum cylinder of different length. A sinusoidal pulse was transmitted across the floor of the pond and the reflected signals were recorded over a period of time sampled at uniform rate. The data obtained contain waveforms reflected from the entire area of the pond floor. Waveforms corresponding to objects are extracted and preprocessed using the algorithm described in [25, Chaps. 3, 4], which is an improved version of the algorithm presented in [26]. This yields ensembles of waveforms per object, and each ensemble consists of rectangular blocks of waveforms.

Our goal is to identify objects according to their material compositions regardless of their shapes. We name the sphere and the cylinder in Experiment j as S_j and C_j , for $j = 1, 2, 3$. Sphere S_1 was filled with air, so we categorize it as one class with label **IA** for iron-air. Spheres S_2 and S_3 were filled with silicone oil so we group them into another class with label **IS** for iron-silicone. All three cylinders were of the same diameter and of the same material, so we grouped them into one class with label **AI** for aluminum. We note, however, that the physical length of C_1 and that of C_2 were the same while that of C_3 was slightly shorter than the half of C_1 and C_2 .

These waveform data are of extremely high dimension: each data point \mathbf{x}_i is a rectangular block of waveforms (i.e., 2D array) of size 17 (cross-range samples) by 600 (time samples), i.e., $d = 17 \times 600$; see Fig. 11.4 for some examples. As for the number of ensembles, we have six ensembles corresponding to these six objects, S_j , C_j , $j = 1, 2, 3$. The number of data points (blocks of waveforms) contained in each ensemble was 8, 8, 16 for $S1$, $S2$, $S3$, respectively, while that of the three cylinders was 32 each. We set aside one ensemble of waveforms (corresponding to one object) as a test ensemble and trained our algorithm on the remaining five sets. For example, when we use the $S1$ ensemble as a test ensemble, we have, using our notation, $\mathcal{X} = \bigcup_{i=1}^5 \mathcal{X}^i$, $\mathcal{Y} = \mathcal{Y}^1$, i.e., $M = 5$ and $N = 1$ while the number of data points (each of which is of size 17×600) in each ensemble is, $m_1 = 8$, $m_2 = 16$, $m_3 = m_4 = m_5 = 32$, and $n_1 = 8$. Then we applied the steps in Algorithm 11.4.1 to classify the test object. We cycled through all six objects, that is, we repeated the classification process six times. The classification results for all six runs are shown in Table 11.1. Using EMD coupled with nonlinear dimensionality reduction methods (DM, LE_{rw} , LE_{sym}), we consistently and correctly classify all three cylinders as objects of class **AI** and the spheres $S2$ and $S3$ as objects of class **IS**. Moreover, the mistake of labeling the sphere $S1$ as **AI** is also consistent. Note that this error is expected since the class **IA** contains only one member ensemble $S1$. We have no training data for this class when the sphere $S1$ is left out as test data. Furthermore, note that having $S1$ as part of the training data does not confuse the classification of the spheres $S2$ and $S3$ when one of the nonlinear dimensionality reduction methods is applied. This is not the case when PCA is used for dimensionality reduction. We will discuss more on this phenomenon in Sect. 11.6 below.

Classification of the objects using HD is less consistent among the different dimensionality reduction/embedding methods. The main reason for this is that HD is highly sensitive to outliers. For a closer look, let us examine the distribution of the points embedded by the Laplacian eigenmap (the random walk version) into the lower-dimensional space. Figure 11.2 shows three sets of points embedded into the first three coordinates of LE_{rw} computed from the training data consisting of all three cylinders and the spheres $S1$ and $S3$. The sphere $S2$ is first left out as test

Table 11.1 Identification of underwater objects

Object		C1	C2	C3	S1	S2	S3
True label		AI	AI	AI	IA	IS	IS
PCA	EMD	AI	AI	AI	IS	IS	IA
	HD	AI	AI	AI	IS	IS	IA
LE_{rw}	EMD	AI	AI	AI	AI	IS	IS
	HD	AI	AI	AI	AI	AI	IS
LE_{sym}	EMD	AI	AI	AI	AI	IS	IS
	HD	AI	AI	AI	AI	IS	IS
DM	EMD	AI	AI	AI	AI	IS	IS
	HD	AI	AI	AI	AI	IS	IS




Fig. 11.2 Three underwater objects in the first three LE_{rw} coordinates. $C3$ (crosses), $S3$ (triangles), unlabeled test object $S2$ (circles). Stars indicate the cluster centers. The Laplacian eigenmap in this case was computed from the training data consisting of all three cylinders and the spheres $S1$ and $S3$

Table 11.2 EMD and HD values in the LE_{rw} coordinates between sphere object $S2$ and all other objects. ($k = 13$ for this example)

Object	C1	C2	C3	S1	S3
EMD	0.0070	0.0064	0.0057	0.0085	0.0053
HD	0.1917	0.2374	0.1237	0.1500	0.1684

data then embedded into the same reduced embedding space via the GHME scheme. In this figure, the crosses correspond to cylinder $C3$ (class **AI**), the triangles correspond to sphere $S3$ (class **IS**), and the circles correspond to the unlabeled test object $S2$ (true label is **IS**). The stars are the cluster centers, i.e., the representatives, in the signature of each object. We see that the circles (corresponding to object $S2$) are on average close to the triangles (corresponding to object $S3$), but because of the points along the long tail of $S3$, the HD between $S2$ and $S3$ turns out to be larger than that between $S2$ and $C3$. The actual EMD and HD values are shown in Table 11.2. Note that the dimension s of the reduced embedding space is actually 12 not 3 in this case. We can see from Table 11.2 that the smallest EMD value is 0.0053 corresponding to $S3$ and the smallest HD value is 0.1237 corresponding to $C3$. Thus, EMD correctly labels object $S2$ as **IS**, but HD mislabels $S2$ as **AI**.

11.5.2 Lip Reading Experiment

In this section, we present our results on a simplified version of the lip-reading problem to illustrate how our proposed algorithm can be applied in practice. The objective of lip reading is to train a machine to automatically recognize the spoken words from the movements of the lips captured on silent video segments (no sound is involved). Much research effort has been devoted to this area. Many published algorithms involve sophisticated feature selection. In this example, we simply perform dimensionality reduction on the sequences of images (i.e., video frames). We do not extract any particular lip features from the images unlike those in [31,40] and many other publications. Furthermore, the lips data we use were collected from one speaker. It may be necessary to extract and use more sophisticated features when multiple speakers, i.e., more variations in the lips, are involved.

We recorded a subject speaking the first five digits (“one”, ..., “five”) ten times using a Nikon Coolpix digital camera sampling at a rate of 60 frames per second. We then extracted the video frames from each movie clip and did the following simple preprocessing. First, we converted the images from color to gray scales ranging from 0 to 255. Then we cropped each image to a 55×70 pixels window around the lips to compensate for translations. (The speaker’s nose was marked with a color marker to facilitate automatic cropping of the video frames). Figure 11.3 shows subsets of video frames of two such movie clips. For each spoken digit, we randomly selected five video sequences from the available ten such sequences in our collection as training ensembles. This gave us a total of 25 video sequences as the training ensembles and another 25 for test ensembles. Hence, using our notation, our experimental setting for each run is: $\mathcal{X} = \bigcup_{i=1}^{25} \mathcal{X}^i$, $\mathcal{Y} = \bigcup_{j=1}^{25} \mathcal{Y}^j$, i.e., $M = N = 25$, while $30 \leq m_i, n_j \leq 63$ (i.e., the number of video frames in each sequence varies between 30 and 63), and the dimension of each data point (i.e., a video frame) is $d = 55 \times 70$.

We applied Algorithm 11.4.1 to classify the test video sequences. We repeated the whole process 100 times. The total misclassification rates (averaging over 100 experimental trials) are shown in Table 11.3. Again, we see that using EMD gives smaller recognition errors than using HD except the LE_{rw} case where the results tied. Table 11.3 shows high classification errors for the Laplacian eigenmaps,



Fig. 11.3 Examples of the preprocessed video frames used in our lip-reading experiments. Both the top row and the bottom row are the subsets of the video frames speaking the digit “one”. Note that the bottom row corresponds to the situation when the subject spoke “one” while smiling

Table 11.3 Lip-reading total recognition errors (averaged over 100 experimental trials)

PCA EMD	PCA HD	LE _{rw} EMD	LE _{rw} HD	LE _{sym} EMD	LE _{sym} HD	DM EMD	DM HD
5.3%	9.4%	36.1%	36.1%	26.0%	27.6%	24.1%	25.2%

in particularly, LE_{rw} while those of DM are at least more than 10% better. Compared to these nonlinear dimensionality reduction methods, however, PCA, i.e., the linear method, worked much better for this dataset, i.e., 5.3% with EMD and 9.4% with HD.

11.6 Discussion and Conclusion

We have seen that PCA performed better than the local nonlinear methods in the lip-reading experiment whereas the situation was opposite in the underwater objects classification experiment. The reason is because the characterizing patterns in the lip-reading problem are “global” whereas those in the underwater-objects problem are “local”. To explain this in more details, suppose we view a video sequence by eyes. In order for us to determine what the spoken word actually is, we need to see how the lips move throughout the entire video sequence. We do not care so much that the shape of the lips in each individual video frame is slightly different from what we have seen in the past. As long as the dynamics (or total movements) of the lips in this video sequence are similar to what we remember, we would be able to recognize the spoken word. In other words, the *trajectory* of the video frames in the embedded space is decisive. Recognizing and discriminating such trajectories is a “global” pattern recognition problem. We believe that is the reason why PCA in the lip-reading experiment outperformed the diffusion maps and Laplacian eigenmaps.

On the other hand, the sonar waveform blocks of the same class are quite homogeneous. To illustrate this, we display in Fig. 11.4 some selected samples of waveforms reflected from the spheres S_1 and S_3 in the underwater-objects experiment. Recall that each rectangular block is viewed as one data point. We can see clearly that the blocks belonging to sphere S_1 are quite different from those from sphere S_3 while our eyes can barely discern the differences between the blocks within the same class. Local nonlinear methods can map each class to the tighter “localized” clusters and enhance the between-class differences compared to PCA. This is the reason why these local nonlinear methods outperformed PCA for this example.

In conclusion, we have proposed an algorithm for classifying objects that are characterized or described by ensembles of signals using a low-dimensional embedding and the Earth Mover’s Distance (EMD). Our algorithm sets up the framework for application of EMD to such classification problems. We have shown that EMD is more robust to noise and hence more appropriate for discrimination of ensembles than the Hausdorff distance.

We have provided two examples of practical applications for our proposed algorithm. The lip-reading application is of global-pattern nature, therefore




Fig. 11.4 Selected waveforms corresponding to the sphere *S1* (top row) and sphere *S3* (bottom row) in the underwater object experiment. The horizontal axis represents time (600 time samples) whereas the vertical axis indicates the 17 cross-range coordinates

dimensionality reduction by PCA proved successful. On the other hand, the classification of underwater objects is of local-pattern nature, thus dimensionality reduction by nonlinear local methods such as diffusion maps and Laplacian eigenmaps gave better results than PCA.

We did not incorporate the discriminant information during the training stage explicitly in this article. For example, we have not compared the performance of our proposed method with the other typical classification strategy, i.e., extracting discriminant features via e.g., Local Discriminant Basis [36,37] followed by setting up a classifier, e.g., Linear Discriminant Analysis [17, Chap. 4], Support Vector Machines [17, Chap. 12], etc., on the extracted feature vectors. We plan to investigate the embedding techniques explicitly incorporating the discriminant information during the training stage. Some of our attempts along this direction can be found in [25].

Acknowledgements This work was partially supported by the ONR grants N00014-06-1-0615, N00014-07-1-0166, N00014-09-1-0041, N00014-09-1-0318, the NSF grant DMS-0410406, and the NSF VIGRE grants DMS-0135345, DMS-0636297. A preliminary version of a subset of the material in this article was presented at the SPIE Wavelets XII Conference, in San Diego, in August 2007 [22]. We thank Dr. Quyen Huynh and Dr. Joe Lopes of NSWC-PC for providing the experimental sonar data. We have used the SPECTRAL Toolbox version 0.1 distributed on the web by Dr. Guido Sanguinetti and Dr. Jonathan Laidler to compute Elongated Kmeans. Finally, we would like to thank Dr. Bradley Marchand for his help in processing the sonar data used in the numerical experiments and Mr. Julien van Hout for his help in numerical experiments on underwater object classification.

References

1. Basseville, M.: Distance measures for signal processing and pattern recognition. *Signal Processing* **18**(4), 349–369 (1989)
2. Belkin, M., Niyogi, P.: Laplacian eigenmaps and spectral techniques for embedding and clustering. In: T.K. Leen, T.G. Dietterich, V. Tresp (eds.) *Advances in Neural Information Processing Systems*, vol. 13, pp. 585–591. The MIT Press, Cambridge, MA (2001)

3. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* **15**(6), 1373–1396 (2003)
4. Belkin, M., Niyogi, P.: Towards a theoretical foundation for Laplacian-based manifold methods. *J. Comput. Syst. Sci.* **74**(8), 1289–1308 (2008)
5. Bengio, Y., Paiement, J.F., Vincent, P., Delalleau, O., Roux, N.L., Ouimet, M.: Out-of-sample extensions for LLE, Isomap, MDS, Eigenmaps, and spectral clustering. In: S. Thrun, L.K. Saul, B. Schölkopf (eds.) *Advances in Neural Information Processing Systems*, vol. 16, pp. 177–184. The MIT Press, Cambridge, MA (2004)
6. Bıyıkoglu, T., Leydold, J., Stadler, P.F.: Laplacian Eigenvectors of Graphs, *Lecture Notes in Mathematics*, vol. 1915. Springer-Verlag, New York (2007)
7. Borg, I., Groenen, P.J.F.: *Modern Multidimensional Scaling: Theory and Applications*, 2nd edn. Springer, New York (2005)
8. Chung, F.R.K.: Spectral Graph Theory. No. 92 in CBMS Regional Conference Series in Mathematics. Amer. Math. Soc., Providence, RI (1997)
9. Coifman, R.R., Lafon, S.: Diffusion maps. *Applied and Computational Harmonic Analysis* **21**(1), 5–30 (2006)
10. Coifman, R.R., Lafon, S.: Geometric harmonics. *Applied and Computational Harmonic Analysis* **21**(1), 32–52 (2006)
11. Cover, T.M., Hart, P.: Nearest neighbor pattern classification. *IEEE Trans. Inform. Theory* **IT-13**, 21–27 (1967)
12. Donoho, D.L., Grimes, C.: Hessian eigenmaps: Locally linear embedding technique techniques for high-dimensional data. *P. Natl. Acad. Sci. USA* **100**(10), 5591–5596 (2003)
13. Evans, L.C.: Partial differential equations and Monge-Kantorovich mass transfer. Tech. rep., Dept. Math., Univ. California, Berkeley (2001)
14. Fowlkes, C., Belongie, S., Chung, F., Malik, J.: Spectral grouping using the Nyström method. *IEEE Trans. Pattern Anal. Machine Intell.* **26**(2), 214–225 (2004)
15. Gower, J.C.: Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika* **53**(3/4), 325–338 (1966)
16. Haker, S., Zhu, L., Tannenbaum, A., Angenent, S.: Optimal mass transport for registration and warping. *Intern. J. Comput. Vision* **60**(3), 225–240 (2004)
17. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd edn. Springer-Verlag (2009)
18. Kirby, M., Sirovich, L.: Application of the Karhunen–Loève procedure for the characterization of human faces. *IEEE Trans. Pattern Anal. Machine Intell.* **12**(1), 103–108 (1990)
19. Lafon, S., Keller, Y., Coifman, R.R.: Data fusion and multicue data matching by diffusion maps. *IEEE Trans. Pattern Anal. Machine Intell.* **28**(11), 1784–1797 (2006)
20. Lafon, S., Lee, A.B.: Diffusion maps and coarse-graining: A unified framework for dimensionality reduction, graph partitioning and data set parameterization. *IEEE Trans. Pattern Anal. Machine Intell.* **28**(9), 1393–1403 (2006)
21. Lafon, S.S.: Diffusion maps and geometric harmonics. Ph.D. thesis, Dept. Math., Yale Univ. (2004). Downloadable from <http://www.math.yale.edu/~sl349>
22. Lieu, L., Saito, N.: Automated discrimination of shapes in high dimensions. In: D. Van De Ville, V.K. Goyal, M. Papadakis (eds.) *Wavelets XII*, Proc. SPIE 6701 (2007). Paper # 67011V
23. von Luxburg, U.: A tutorial on spectral clustering. *Stat. Comput.* **17**(4), 395–416 (2007)
24. van der Maaten, L.J.P., Postma, E.O., van den Herik, H.J.: Dimensionality reduction: A comparative review. Technical Report TiCC-TR 2009-005, Tilburg Centre for Creative Computing, Tilburg Univ. (2009)
25. Marchand, B.: Local signal analysis for classification. Ph.D. thesis, Dept. Math., Univ. California, Davis (2010)
26. Marchand, B., Saito, N., Xiao, H.: Classification of objects in synthetic aperture sonar images. In: Proc. 14th IEEE Workshop on Statistical Signal Processing, pp. 433–437. IEEE (2007)
27. Mardia, K.V., Kent, J.T., Bibby, J.M.: *Multivariate Analysis*. Academic Press, San Diego, CA (1979)
28. Murase, H., Nayar, S.K.: Visual learning and recognition of 3d objects from appearance. *Intern. J. Comput. Vision* **14**(1), 5–24 (1995)

29. Nesbitt, C.L., Lopes, J.L.: Subcritical detection of an elongated target buried under a rippled interface. In: *Oceans '04, MTS/IEEE Techno-Ocean '04*, vol. 4, pp. 1945–1952 (2004)
30. Ng, A.Y., Jordan, M.I., Weiss, Y.: On spectral clustering: Analysis and an algorithm. In: T. Dietterich, S. Becker, Z. Ghahramani (eds.) *Advances in Neural Information Processing Systems*, vol. 14, pp. 849–856. The MIT Press, Cambridge, MA (2002)
31. Patterson, E.K., Gurbuz, S., Tufekci, Z., Gowdy, J.N.: Moving-talker, speaker-independent feature study, and baseline results using the CUAVE multimodal speech corpus. *EURASIP J. Appl. Signal Process.* **11**, 1189–1201 (2002)
32. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**(5500), 2323–2326 (2000)
33. Rubner, Y., Tomasi, C.: *Perceptual Metrics for Image Database Navigation*. Kluwer Academic Publishers, Boston, MA (1999)
34. Rubner, Y., Tomasi, C., Guibas, L.J.: The Earth Mover's Distance as a metric for image retrieval. *Intern. J. Comput. Vision* **40**(2), 99–121 (2000)
35. Saito, N.: Image approximation and modeling via least statistically dependent bases. *Pattern Recognition* **34**, 1765–1784 (2001)
36. Saito, N., Coifman, R.R.: Local discriminant bases and their applications. *J. Math. Imaging Vis.* **5**(4), 337–358 (1995). Invited paper
37. Saito, N., Coifman, R.R., Geshwind, F.B., Warner, F.: Discriminant feature extraction using empirical probability density estimation and a local basis library. *Pattern Recognition* **35**(12), 2841–2852 (2002)
38. Saito, N., Woei, E.: Analysis of neuronal dendrite patterns using eigenvalues of graph Laplacians. *JSIAM Letters* **1**, 13–16 (2009). Invited paper
39. Sanguinetti, G., Laidler, J., Lawrence, N.D.: Automatic determination of the number of clusters using spectral algorithms. In: *Proc. 15th IEEE Workshop on Machine Learning for Signal Processing*, pp. 55–60 (2005)
40. Zhang, X., Mersereau, R.M.: Lip feature extraction towards an automatic speechreading system. In: *Proc. 2000 International Conference on Image Processing*, vol. 3, pp. 226–229. IEEE (2000)
41. Zhang, Z., Zha, H.: Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM J. Sci. Comput.* **26**(1), 313–338 (2005)
42. Zhou, S.K., Chellappa, R.: From sample similarity to ensemble similarity – Probabilistic distance measures in reproducing kernel Hilbert space. *IEEE Trans. Pattern Anal. Machine Intell.* **28**(6), 917–929 (2006)

Part III

Applications of Wavelets

Chapter 12

Wavelets on Manifolds and Statistical Applications to Cosmology

Daryl Geller and Azita Mayeli

Abstract We outline how many of the pioneering ideas that were so effective in developing wavelet theory on the real line, can be adapted to the manifold setting. In this setting, however, arguments using Fourier series are replaced by methods from modern harmonic analysis, pseudodifferential operators and PDE. We explain how to construct nearly tight frames on any smooth, compact Riemannian manifold, which are well-localized both in frequency (as measured by the Laplace–Beltrami operator) and in space. These frames can be used to characterize Besov spaces on the manifold for the full range of indices, in analogy to the Frazier–Jawerth result on the real line. We explain how our methods can be extended beyond the study of functions, to the wavelet analysis of sections of a particular line bundle on the sphere, which is important for the analysis of the polarization of CMB (cosmic microwave background radiation). The wavelet approach to CMB has been advocated by many people, including our frequent collaborators, the statistician Domenico Marinucci, and the physicist Frode Hansen, who earlier used spherical needlets to study CMB temperature.

12.1 Introduction

The original architects of wavelet theory worked on the real line, and ingeniously constructed frames and orthonormal wavelet bases, largely through clever use of the interplay between Fourier series and the Fourier transform. Because of natural applications to time–frequency analysis, it was very natural to work on the real line.

Our main goal in this article is to outline how many of the pioneering ideas that were so effective on the real line, can be adapted to the manifold setting. In this setting, however, as we shall explain in the next section, arguments using Fourier series are replaced by methods from modern harmonic analysis, pseudodifferential operators and PDE.

D. Geller (deceased)

Mathematics Department, Stony Brook University, Stony Brook NY, 11794, USA

Although there are many problems on curved surfaces for which one can work very locally (so that geometrical considerations related to the Riemannian metric become irrelevant) there are also problems where such a localization is entirely inappropriate. Specifically, in the later sections of this article, we shall explore some applications to cosmology, in particular the statistical analysis of the CMB (cosmic microwave background radiation), which has been called the “Rosetta Stone” of the Big Bang. In this problem, the desire to exploit the natural assumption that “physics is the same in all directions” leads inexorably to truly spherical analysis. Spherical harmonics are a primary tool of all the major research teams [31, 34]; but the use of spherical wavelets is being advocated by some prominent astrostatisticians and astrophysicists, including Domenico Marinucci and Frode Hansen, who collaborated with us on much of the work we will describe in Sects. 12.6 and 12.7. A major reason for the interest in wavelets in CMB is that a large portion of the sky is not observed, because of interference from the Milky Way and other sources. These ideas have been exploited for a number of years to study CMB temperature [1, 3–6, 20, 24, 27–29, 37, 38]. As we shall explain, the authors’ main contribution in this regard [11–13, 17] is that one can study CMB *polarization* by wavelet methods as well. Polarization must be regarded as a section of a line bundle, and is not an ordinary scalar function like temperature. We have also proposed [15, 16, 26, 33] a new spherical wavelet frame (Mexican needlets) for the analysis of CMB temperature, which appears [33] to have numerical advantages over earlier methods.

Spherical wavelets have also had numerous applications to geoscience, especially by Freedman and collaborators [9, 10]. This includes the science of GPS and satellite-to-satellite communication, both of which require one to work on a large portion of a sphere, so that curvature considerations become essential. Spherical harmonic coefficients will all change if local information changes, which makes them computationally inferior to wavelet coefficients. Some of the wavelets we consider on the sphere are quite similar to those considered by Freedman et al. The crucial new ingredient is that we now have much more information about such wavelets in space (as opposed to frequency, where they are quite explicit), not only on the sphere, but on general smooth compact Riemannian manifolds. Using this information, we obtain nearly tight frames by a spatial discretization process (Freedman et al worked in the much easier continuous wavelet setting for L^2). Moreover, we obtain a general wavelet-based analysis of function spaces far beyond L^2 , analogous to that carried out by Frazier and Jawerth [8] on the real line.

12.2 A New Look at Wavelets on the Real Line

In order to illustrate our general philosophy for constructing wavelets on manifolds, we re-examine the very familiar situation on the real line. Specifically, we address the question of how one can discretize continuous wavelets into nearly tight frames. Our discussion will be fairly standard, but with one important change – we will not use Fourier series arguments. This change makes all the difference, since it provides a path through which one can generalize the arguments to manifolds.

Say $0 \neq f \in \mathcal{S}(\mathbb{R}^+)$ (the space of restrictions to \mathbb{R}^+ of functions in $\mathcal{S}(\mathbb{R})$), but that $f(0) = 0$. One then has the *Calderón reproducing formula* : if $c \in (0, \infty)$ is defined by

$$c = \int_0^\infty |f(t)|^2 \frac{dt}{t},$$

then for all $s > 0$,

$$\int_0^\infty |f(ts)|^2 \frac{dt}{t} = c < \infty. \quad (12.1)$$

Discretizing this formula, one obtains a special form of *Daubechies' condition* [7]: if $a > 1$ is sufficiently close to 1, then for all $s > 0$,

$$0 < A_a \leq \sum_{j=-\infty}^{\infty} |f(a^{2j}s)|^2 \leq B_a < \infty. \quad (12.2)$$

Moreover, by use of the standard error estimate for the midpoint rule in evaluating integrals, one finds [14]

$$A_a = \frac{c}{2|\log a|} (1 - O(|(a-1)^2(\log|a-1|)|)), \quad (12.3)$$

$$B_a = \frac{c}{2|\log a|} (1 + O(|(a-1)^2(\log|a-1|)|)). \quad (12.4)$$

so that, as $a \rightarrow 1^+$, $B_a/A_a \rightarrow 1$ nearly quadratically. Thus, one would expect that, in practical examples, B_a/A_a would be very close to 1 even if a is at some distance from one. In fact, in the important special case $f(s) = se^{-s}$, $a = 2^{1/3}$, Daubechies calculated [7] that $B_a/A_a = 1.0000$ to four digits. As we shall see in a moment, this is important for the construction of practical nearly tight frames.

Next, note that the even function

$$h(\xi) = f(\xi^2)$$

is in $\mathcal{S}(\mathbb{R})$ and satisfies

$$\int_0^\infty |h(t\xi)|^2 \frac{dt}{t} = \frac{c}{2} < \infty, \quad (12.5)$$

and

$$0 < A_a \leq \sum_{j=-\infty}^{\infty} |h(a^j\xi)|^2 \leq B_a < \infty, \quad (12.6)$$

both for $\xi \neq 0$. For instance, if $f(s) = se^{-s}$, then $h(\xi) = \xi^2 e^{-\xi^2}$. If

$$\psi = \check{h}$$

then ψ is (a constant multiple of) the *Mexican hat function*, the second derivative of a Gaussian. If $F \in L^2(\mathbb{R})$, we can multiply both sides of (12.5) by $|F(\xi)|^2$, integrate with respect to ξ , and use Plancherel to find

$$\int_0^\infty \|F * \psi_t\|_2^2 \frac{dt}{t} = \frac{c}{2} \|F\|_2^2. \quad (12.7)$$

if $\psi_t(x) = (1/t)\psi(x/t)$. The same procedure, using (12.6), but now also using a spatial discretization argument, yields that for b sufficiently small,

$$A\|F\|_2^2 \leq \sum |\langle F, \phi_{jk} \rangle|_2^2 \leq B\|F\|_2^2, \quad (12.8)$$

$$\phi_{jk}(x) = \sqrt{ba^j} \bar{\psi}_{a^j}(x - bka^j) \quad (12.9)$$

$$\frac{B}{A} \sim \frac{B_a}{A_a} \sim 1. \quad (12.10)$$

if a is reasonably close to 1. Thus, for sufficiently small b , the ϕ_{jk} are a *nearly tight* frame.

The key point we need to highlight here is the spatial discretization argument which enables one to obtain (12.8). In performing the spatial discretization, Daubechies (in Chap. 3 of her book) obtains estimates for the errors by using Fourier series. But one can also [14–16] use the $T(1)$ theorem from harmonic analysis, which can be used not just to tell you when you have a Calderón–Zygmund operator, but also to estimate its norm on L^2 .

Key to this is evidently the operator of convolution with ψ_t .

When using the $T(1)$ theorem, it is crucial that ψ_t and its derivatives satisfy the estimates that an approximate identity does. (But note that the integral of each ψ_t is zero, since its Fourier transform vanishes at the origin.) We return to the crucial estimates satisfied by the ψ_t in Sect. 12.4 below.

12.3 Wavelets on Manifolds: The Set-Up

Remarkably, one can generalize this procedure to *many* different situations. In fact, if T is a positive self-adjoint operator on a Hilbert space \mathcal{H} , from

$$\int_0^\infty |f(ts)|^2 \frac{dt}{t} = c < \infty \quad (s > 0)$$

$$0 < A_a \leq \sum_{j=-\infty}^\infty |f(a^{2j}s)|^2 \leq B_a < \infty \quad (s > 0)$$

we find [15] that if P is the projection onto the null space of T ,

$$\int_0^\infty |f|^2(tT) \frac{dt}{t} = c(I - P), \quad (12.11)$$

$$A_a(I - P) \leq \sum_{j=-\infty}^\infty |f|^2(a^{2j}T) \leq B_a(I - P). \quad (12.12)$$

Here, the integral and the sum converge strongly. (For self-adjoint operators R, S on \mathcal{H} , as usual $R \leq S$ means $\langle Rx, x \rangle \leq \langle Sx, x \rangle$ for all $x \in \mathcal{H}$.)

To prove these facts we use the *spectral theorem*.

On the real line the relevant operator T is $-d^2/dx^2$. Then if $F \in L^2(\mathbb{R})$,

$$(f(t^2T)F)^\wedge(\xi) = f(t^2\xi^2)\hat{F}(\xi) = h(t\xi)\hat{F}(\xi) = (\psi_t * F)^\wedge(\xi),$$

so that $f(t^2T)$ is just convolution with ψ_t , the key operator used above. We could write

$$[f(t^2T)F](x) = \int_0^\infty K_t(x, y)F(y)dy$$

with

$$K_t(x, y) = \psi_t(x - y).$$

We again emphasize that the estimates satisfied by K_t and its derivatives are crucial to the analysis.

Now instead take $T = \Delta$, the Laplace–Beltrami operator ($= -d^*d$) on a smooth compact (connected) Riemannian manifold \mathbf{M} . Write

$$[f(t^2T)F](x) = \int_{\mathbf{M}} K_t(x, y)F(y)d\mu(y)$$

for a suitable K_t . (Here μ is Riemannian measure on \mathbf{M} .)

We show that K_t satisfies similar estimates as in the case of the real line. From this and similar arguments to those we gave on \mathbb{R} , but now using

$$A_a(I - P) \leq \sum_{j=-\infty}^\infty |f|^2(a^{2j}T) \leq B_a(I - P)$$

in place of the Daubechies condition, we show:

Theorem 12.1. [16] *Pick $b > 0$, and for each j , write \mathbf{M} as a disjoint union of measurable sets $E_{j,k}$ with diameter at most ba^j . Assume also the $E_{j,k}$ are also not too small, in the sense that, for some $c_0 > 0$ and some J , $\mu(E_{j,k}) \geq c_0(ba^j)^n$ whenever $j < J$. Take $x_{j,k} \in E_{j,k}$.*

Let

$$\phi_{j,k}(x) = \sqrt{\mu(E_{j,k})} \overline{K_{a^j}}(x_{j,k}, x). \quad (12.13)$$

Then, if b is sufficiently small, $\{\phi_{j,k}\}$ is a nearly tight frame for $(I - P)L^2(\mathbf{M})$.

Note that these $\phi_{j,k}$ are entirely analogous to (12.9) on the real line. We make the following remarks:

1. The frame elements can be written in the form

$$\phi_{j,k}(x) = \sqrt{\mu(E_{j,k})} \overline{K_{a^j}}(x_{j,k}, x) = \sqrt{\mu(E_{j,k})} \sum_l \bar{f}(a^j \lambda_l) u_l(x_{j,k}) u_l(x),$$

where the u_l are an orthonormal basis of real-valued eigenvectors for Δ , and the λ_l are the corresponding eigenvalues. Such expressions have arisen in numerous contexts.

2. We can characterize [18] the Besov spaces on \mathbf{M} , for the full range of indices, through a knowledge of the size of frame coefficients. (Depending on the indices, one might need f to vanish to some order greater than 1 at the origin.)
3. Let us say that f is *band-limited* if $\text{supp } f \subseteq [c, d]$ for some $0 < c < d < \infty$. In our theory, we *do not need to require* that f is band-limited. Band-limited f and non-band-limited f each have their own advantages. For particular Band-limited f , one can replace the Daubechies condition

$$0 < A_a \leq \sum_{j=-\infty}^{\infty} |f(a^{2j}s)|^2 \leq B_a < \infty.$$

by

$$\sum_{j=-\infty}^{\infty} |f(a^{2j}s)|^2 \equiv 1$$

for all $s > 0$. Then one can often arrange to discretize the spatial integrals *exactly* by using *cubature points* to obtain frames that are *tight*, not just nearly tight.

This has been done in a number of contexts: early on, on the real line, by Daubechies, Meyer, and Frazier and Jawerth, and more recently on the sphere by Narcowich, Petrushev and Ward [28, 29]. This was generalized to compact homogeneous manifolds (smooth compact manifolds, for which a Lie group acts transitively through diffeomorphisms) by Geller and Pesenson in [19]; in particular one can obtain band-limited, smooth, localized tight frames on compact Lie groups.

On the other hand, if f is not required to be band-limited, one can arrange [15] for the frame element $\phi_{j,k}(x)$ to be compactly supported within a ball of radius a^j (the appropriate scale) (by taking appropriate f). Indeed, recall $h(\xi) = f(\xi^2)$ is even. If $\text{supp } \hat{h} \subseteq (-1, 1)$ then

$$f(t^2 \Delta) = h(t\sqrt{\Delta}) = c \int_{-1}^1 \hat{h}(\xi) \cos(\xi t\sqrt{\Delta}) d\xi \quad (12.14)$$

has kernel K_t supported in $\{(x, y) : d(x, y) \leq t\}$, by the finite speed of propagation property of the wave equation. This yields the desired support condition, by (12.13).




Fig. 12.1 Mexican hat wavelet on S^2 for $t = 2$ (left), $t = 1/2$ (middle), $t = 1/8$ (right)

The special non-band-limited case of $f(s) = se^{-s}$ (the “Mexican hat”) has a particular appeal. In this situation, one conjecturally expects Gaussian decay of the frame elements $\phi_{j,k}$, and minimal oscillation of the $\phi_{j,k}$. (In contrast, one expects the frame elements in the band-limited case to look like $\sin x/x$, so that calculations on the “space” side, (as opposed to the “frequency” side) would not be as numerically stable in the band-limited situation.) Moreover, one expects to be able to compute the $\phi_{j,k}(x)$ very accurately, even at small scales. On the torus, these conjectures can be verified [15] by use of the Poisson summation formula. On the sphere S^2 , the theory appears within reach [15]. See Fig. 12.1 for pictures of (a constant multiple of) the Mexican hat wavelet $K_t(\mathbf{N}, x)$ on S^2 (here $\mathbf{N} = (0, 0, 1)$ is the North pole, and $f(s) = se^{-s}$), for $t = 2, 1/2, 1/8$. In this figure, the horizontal axis is the usual angle θ in spherical coordinates, which we have allowed to range from $-\pi$ to π . The wavelet is a function of θ only.

On the sphere, when $f(s) = se^{-s}$ is used, we have called the frame $\{\phi_{j,k}\}$ *Mexican needles*. As we will discuss at the end of the article, these appear to have numerical advantages over earlier wavelets used in cosmology.

12.4 Functions of Elliptic Operators

To do the analysis required to make it possible to use the $T(1)$ theorem, or characterize Besov spaces, we need to show that $K_t(x, y)$ (on \mathbf{M}) satisfies estimates, similar to those satisfied by $\psi_t(x - y)$ (on \mathbb{R}).

On \mathbb{R} :

$$|\partial_x^\alpha \partial_y^\beta \psi_t(x - y)| \leq C_{\alpha\beta N} \frac{t^{-1-|\alpha|-|\beta|}}{(1 + |x - y|/t)^N}$$

since $\psi_t(x - y) = t^{-1} \psi((x - y)/t)$, and $\psi \in \mathcal{S}$.

On \mathbf{M} we analogously have [15]:

For every pair of C^∞ differential operators X (in x) and Y (in y) on \mathbf{M} , and for every integer $N \geq 0$,

$$|XYK_t(x, y)| \leq C_{N,X,Y} \frac{t^{-n-\deg X-\deg Y}}{\left(1 + \frac{d(x,y)}{t}\right)^N}. \quad (12.15)$$

(Here, d is the geodesic distance. Recall that K_t is the kernel of $f(t^2\Delta)$ ($f \in \mathcal{S}(\mathbb{R}^+)$). We have been assuming that $f(0) = 0$, but the estimate holds without this assumption, provided $|t| < 1$.)

Since functions of elliptic operators on manifolds have been intensively studied, it is perhaps surprising that (12.15) was not previously known. Let us indicate why, in earlier work on manifolds, researchers did not need the full force of (12.15).

Let us look at a basic example. On the circle, where $\Delta = -d^2/d\theta^2$, if $f(u) = e^{-u}$,

$$K_t(\theta, \phi) = \sum_{m=-\infty}^{\infty} e^{-t^2 m^2} e^{im\theta} e^{-im\phi}$$

so that, if $F(\phi) = \sum_m \hat{F}(m) e^{im\phi}$,

$$f(t^2\Delta)F = e^{-t^2\Delta}F = \sum_m e^{-t^2 m^2} \hat{F}(m) e^{im\theta}.$$

$$F(\theta) = \sum_m \hat{F}(m) e^{im\theta},$$

$$\left[e^{-t^2\Delta} F \right](\theta) = \sum_m e^{-t^2 m^2} \hat{F}(m) e^{im\theta}.$$

Traditionally, one sought to prove that $e^{-t^2\Delta}F \rightarrow F$ in various function spaces, as $t \rightarrow 0^+$.

To do this, one would expect to have to show that $e^{-t^2\Delta}F$ is similar to convolving F with an approximate identity. Thus, one might think that researchers would have needed to show (12.15). But actually, it was usually clear (for general f, \mathbf{M}) that $f(t^2\Delta)F \rightarrow F$ for $F \in C^\infty$, which is dense in many function spaces. One then only needed to show that the operators $f(t^2\Delta)$ were *uniformly bounded* as $t \rightarrow 0^+$ on the function spaces, and so much less information was needed.

However, in our theory, we do require the full force of (12.15). Our proof, though rather short, requires something of a change in perspective from the ways in which weaker estimates are proved in the literature. We may summarize our strategy as follows. Each K_t is smooth, but we regard it as the kernel of a pseudodifferential operator in $OPS_{1,0}^m$ for every m . If we can estimate the symbol of this operator, we can estimate the derivatives XYK_t of the kernel K_t . To estimate the symbol, we use the following theorem of Strichartz ([35], or Theorem 1.3, page 296, of [36]):

- If $p(\xi) \in S_1^m(\mathbb{R})$, then $p(\sqrt{\Delta}) \in OPS_{1,0}^m(\mathbf{M})$.

(Here, $S_1^m(\mathbb{R})$ denotes the space of symbols on \mathbb{R} of order m , which depend only on the dual variable ξ .) In fact, we observe that the map $p \rightarrow p(\sqrt{\Delta})$ is continuous from the Fréchet space $S_1^m(\mathbb{R})$ to the Fréchet space $OPS_{1,0}^m(\mathbf{M})$. Thus, to estimate

the symbol of $f(t^2\Delta) = h(t\sqrt{\Delta})$, it is enough to estimate the $S_1^m(\mathbb{R})$ norms of the functions $h_{[t]}(\xi) = h(t\xi)$. This is then elementary.

This summary of our strategy is a little oversimplified. In order to estimate XYK_t in this manner, one would need that h vanish to a high order at the origin (high relative to “ N ” in (12.15)). But one can reduce to this situation by splitting h into two parts, one vanishing to a high order at the origin, and the other with Fourier transform supported in $(-1, 1)$. Because of the finite propagation speed property of the wave equation, and (12.14), the latter part of h leads to a kernel supported in $\{(x, y) : d(x, y) \leq t\}$. Thus, for this part of h , in proving the estimates, one may assume N is zero.

12.5 Analogues of Dilations and Translations

Recall that our wavelets on \mathbf{M} have the form

$$\phi_{j,k}(x) = \sqrt{\mu(E_{j,k})} \overline{K_{aj}}(x_{j,k}, x),$$

On \mathbb{R} , the analogous wavelets (12.9) are obtained by translations and dilations of a fixed function.

On \mathbf{M} , there are no dilations or translations, but:

- K_t is the kernel of $f(t^2\Delta)$, and $t^2\Delta$ would be the Laplace–Beltrami operator if we *dilated* the Riemannian metric (formally $\sum g_{jk} dx_j dx_k$);
- If $\mathbf{M} = G/K_0$ is a homogeneous manifold (G a compact Lie group, K_0 a closed subgroup of G), then for fixed j , the $\overline{K_{aj}}(x_{j,k}, x)$ are all obtained from each other by applying G . For example, on the sphere they are all rotates of each other.

12.6 Spin Wavelets

In this section, we describe a new application, the construction of spin wavelets on the sphere S^2 , with applications to the analysis of CMB (cosmic microwave background radiation).

In 1852, Stokes obtained the following fundamental result in optics. For an electromagnetic field propagating in the z direction,

$$\begin{aligned} E_x &= A_x \cos(kz - \omega t) \\ E_y &= A_y \cos(kz - \omega t + \varphi) \\ E_z &= 0 \end{aligned}$$

the projection onto the (x, y) plane is an ellipse centered at 0, and

$$\begin{aligned} T &= A_x^2 + A_y^2 := I \\ Q &= A_x^2 - A_y^2 = I \cos 2\chi \cos 2\psi \\ U &= 2A_x A_y \cos \phi = I \cos 2\chi \sin 2\psi \\ V &= -2A_x A_y \sin \phi = I \sin 2\chi. \end{aligned}$$

where $|\cot \chi|$ is the ratio of the major and minor axes, and ψ is the angle between the x axis and the major axis. The sign of χ determines if you have right or left polarization. T, Q, U and V are called the *Stokes parameters*. (There are nice pictures of all this on the Wikipedia page on Polarization.)

In CMB (radiation emitted 400,000 years after the Big Bang), these four quantities are all *observable* and *physically measurable* with devices including a polarimeter.

I is the intensity, or temperature. If one time averages over a long time period, one finds that χ averages to zero (so that, on average, there is no left or right polarization); and V averages to zero. But $Q + iU$ averages to $Ie^{2i\psi}$, which must be regarded as a *spin 2 quantity*, for the following reason: There is no preferred choice of the x and y axes. If you rotate them by an angle α , the angle ψ is increased by α , so that the observable quantity $Q + iU = Ie^{2i\psi}$ gets multiplied by $e^{2i\alpha}$.

Let us now explain what is meant by a “spin s quantity”, for s an integer. The physicist’s definition, given by Newman–Penrose [30] is this. A quantity η defined on the sphere has spin weight $s \in \mathbb{Z}$, provided that, whenever a tangent vector m at a point transforms under coordinate change by

$$m' = e^{i\alpha} m,$$

the quantity η , at that point, transforms by

$$\eta' = e^{is\alpha} \eta.$$

Then $Q + iU$ is a spin 2 quantity.

We now explain the mathematical meaning of spin weight s ; for a fuller discussion, see [13]. We let \mathbf{N} be the north pole $(0, 0, 1)$, let \mathbf{S} be the south pole, $(0, 0, -1)$. Let

$$U_I = S^2 \setminus \{\mathbf{N}, \mathbf{S}\}.$$

If $R \in SO(3)$, we define

$$U_R = RU_I.$$

On U_I we use standard spherical coordinates (θ, ϕ) ($0 < \theta < \pi$, $-\pi \leq \phi < \pi$), and analogously, on any U_R we use coordinates (θ_R, ϕ_R) obtained by rotation of the coordinate system on U_I .

At each point p of U_R we let $\rho_R(p)$ be the unit tangent vector at p which is tangent to the circle $\theta_R = \text{constant}$, pointing in the direction of increasing ϕ_R . (This is well-defined since $R\mathbf{N}, R\mathbf{S} \notin U_R$.)

If $p \in U_{R_1} \cap U_{R_2}$, we let

$$\psi_{pR_2R_1} \text{ be the oriented angle from } \rho_{R_1}(p) \text{ to } \rho_{R_2}(p).$$

Note:

- We have changed notation; this ψ has nothing to do with the ψ we used when discussing Stokes parameters.
- At \mathbf{S} , the oriented angle from \mathbf{i} to \mathbf{j} is $\pi/2$; at \mathbf{N} , the oriented angle from \mathbf{j} to \mathbf{i} is $\pi/2$.
- Since R is conformal, the angle $\psi_{pR_2R_1}$ would be clearly be the same if we had instead chosen $\rho_R(p)$ to point in the direction of increasing θ_R , for instance. Thus, ψ_{pRI} measures “the angle by which the tangent plane at p is rotated if the coordinates are rotated by R ”.

Now say $\Omega \subseteq S^2$ is open. Let $C_s^\infty(\Omega) = \{F = (F_R)_{R \in SO(3)} : \text{all } F_R \in C^\infty(U_R \cap \Omega), \text{ and for all } R_1, R_2 \in SO(3) \text{ and all } p \in U_{R_1} \cap U_{R_2} \cap \Omega,$

$$F_{R_2}(p) = e^{is\psi} F_{R_1}(p),$$

where $\psi = \psi_{pR_2R_1}$.

Heuristically, F_R is F_I “looked at after the coordinates have been rotated by R ”; at p , it has been multiplied by $e^{is\psi}$. This is clearly an appropriate way to make the physicists’ concept of spin s precise. Equivalently, $C_s^\infty(\Omega)$ consists of the smooth sections over Ω of the line bundle with transition functions $e^{is\psi_{pR_2R_1}}$ from U_{R_1} to U_{R_2} .

Similarly, we can define $L_s^2(\Omega)$ (the F_R need to be in L^2). There is a well-defined inner product on $L_s^2(\Omega)$, given by $\langle F, G \rangle = \langle F_R, G_R \rangle$; clearly this definition is independent of choice of R .

There is a unitary action of $SO(3)$ on $L_s^2(S^2)$, given by $F \rightarrow F^R$, which is determined by the equation

$$(F^R)_I(p) = F_R(Rp).$$

We think of F^R as a “rotate” of F . Thus we have two important relations: if we “rotate coordinates”, we have

$$F_{R_2}(p) = e^{is\psi} F_{R_1}(p),$$

while if we “rotate F ”, we have

$$(F^R)_I(p) = F_R(Rp).$$

Physicists would say that spin s quantities need to be multiplied by a factor $e^{is\psi}$ when rotated.

Following Newman and Penrose, we define

$$\partial_{sR}F_R = -(\sin \theta_R)^s \left(\frac{\partial}{\partial \theta_R} + \frac{i}{\sin \theta_R} \frac{\partial}{\partial \phi_R} \right) (\sin \theta_R)^{-s} F_R,$$

Then the *spin-raising operator*

$$\partial : C_s^\infty(\Omega) \rightarrow C_{s+1}^\infty(\Omega)$$

given by

$$(\partial F)_R = \partial_{sR} F_R$$

is well-defined. Also, if $\bar{\partial}_{sR} F_R = \overline{\partial_{-s,R} F_R}$, the *spin-lowering operator* $\bar{\partial} : C_s^\infty(\Omega) \rightarrow C_{s-1}^\infty(\Omega)$ given by

$$(\bar{\partial} F)_R = \bar{\partial}_{sR} F_R$$

is well-defined. Also ∂ and $\bar{\partial}$ commute with the actions of $SO(3)$.

If $s \geq 0$, we let

$$\Delta_s = -\bar{\partial}\partial,$$

acting on $C_s^\infty(\Omega)$, while if $s < 0$, we let $\Delta_s = -\partial\bar{\partial}$.

Then Δ_0 is the usual spherical Laplacian (Laplace–Beltrami operator on S^2).

Recall that Δ_0 is elliptic and consequently has an orthonormal basis of eigenfunctions.

When $s = 0$, one lets $\lambda_l = l(l+1)$. For each $l \in \mathbb{N}_0$, there is an orthonormal basis $\{Y_{lm} : -l \leq m \leq l\}$ of eigenfunctions for Δ_0 for the eigenvalue λ_l ; call the eigenspace \mathcal{H}_l . The Y_{lm} are spherical harmonics. One has $L^2(S^2) = \bigoplus_{l=0}^\infty \mathcal{H}_l$. \mathcal{H}_l is analogous to $\langle e^{im\theta}, e^{-im\theta} \rangle$, the eigenspace of the circular Laplacian $-d^2/d\theta^2$ on the circle, for the eigenvalue m^2 .

For general s , we analogously have:

$$L_s^2(S^2) = \bigoplus_{l=|s|}^\infty \mathcal{H}_{ls}.$$

Here, each \mathcal{H}_{ls} is an eigenspace of Δ_s , with an explicit orthonormal basis of *spin spherical harmonics*

$$\{ {}_s Y_{lm} : -l \leq m \leq l \},$$

obtained through application of the spin-raising and spin-lowering operators to the usual spherical harmonics Y_{lm} . In fact, \mathcal{H}_{ls} is the eigenspace of Δ_s for the eigenvalue λ_{ls} , where

$$\lambda_{ls} = \begin{cases} (l-s)(l+s+1) & \text{if } 0 \leq s \leq l, \\ (l+s)(l-s+1) & \text{if } -l \leq s < 0. \end{cases} \quad (12.16)$$

Now fix s . It will be important to note that the map $F \rightarrow F^R$ defines an action of the rotation group acts on *each* eigenspace \mathcal{H}_{ls} of Δ_s . Further, this action is *irreducible*, and the actions for different l are *inequivalent*.

We claim that the theory of Sects. 12.3 and 12.4 has a complete analogue in the spin s situation, even though we are dealing with sections of line bundles and not ordinary functions. Precisely;

set

$$K_{t,R',R}(x,y) = \sum_{l \geq |s|} \sum_{m=-l}^l f(t^2 \lambda_{ls}) {}_s Y_{lmR'}(x) \overline{{}_s Y_{lmR}(y)}$$

then

$$(f(t\Delta_s)F)_{R'}(x) = \int_{S^2} K_{t,R',R}(x,y) f_R(y) d\mu(y).$$

Theorem 12.2. [13, 17] $K_{t,R',R}$ satisfies the usual estimates

$$|XYK_{t,R',R}(x,y)| \leq C_{N,X,Y} \frac{t^{-n-\deg X-\deg Y}}{(1 + \frac{d(x,y)}{t})^N},$$

as long as x stays in a compact subset of $U_{R'}$ and y stays in a compact subset of U_R .

(The last requirement about the location of x and y may be dropped if $X = Y = id$; in general, it poses little problem, since one can always look at K_{t,R'_1,R_1} for other R_1, R'_1 .)

Theorem 12.3. [17] Choose the sets E_{jk} as in Theorem 12.1, pick $x_{jk} \in E_{jk}$, pick R_{jk} with $x_{jk} \in U_{R_{jk}}$, and set

$$\phi_{j,k}(x) = \sqrt{\mu(E_{j,k})} \sum_{l \geq |s|} \sum_{m=-l}^l f(a^{2j} \lambda_{ls}) \overline{{}_s Y_{lmR_{jk}}(x_{jk})} {}_s Y_{lm}(x).$$

so that $\phi_{jk} \in C_s^\infty(S^2)$. Then, for a close enough to 1 and b small enough, the ϕ_{jk} are a nearly tight frame for $(I - P)L_s^2(S^2)$, where $P = \mathcal{H}_{|s|,s}$ is the null space of Δ_s .

(Note: the sum $\sum |\langle F, \phi_{jk} \rangle|^2$ is unaffected by choice of R_{jk} .)

12.7 Applications

We will discuss a new application of these ideas to cosmology in this section. But first we summarize earlier applications here.

12.7.1 Earlier Applications

Equations (12.11) and (12.12) can be used in a wide variety of contexts. Thus, initially, we studied wavelets on stratified Lie groups, like the Heisenberg group, by

using these methods [14, 25]. Meanwhile, Narcowich, Petrushev, and Ward [28, 29] did something similar on the sphere, but only for band-limited f . In particular, they obtained the estimates (12.15), on the sphere, for similar kernels. They obtained exact discretizations and tight frames (*spherical needlets*) by using cubature points; they did not use the $T(1)$ theorem. They also obtained characterizations of Besov and Triebel–Lizorkin spaces on the sphere, using their needlets. As we shall explain in the next section, their needlets have been applied to cosmology.

Our own articles on manifolds [15, 16] first appeared in arXiv in 2007. Petrushev and collaborators (Kyriazis, Xu, \dots), have obtained analogous results in other settings, including the ball [22] and classical orthogonal expansions [21, 32], again by using cubature points.

12.7.2 *New Applications to Cosmology*

CMB (cosmic microwave background radiation) was emitted only about 400,000 years after the Big Bang. Before that, the universe was a hot plasma, and photons had a very short mean free path (they kept bumping into electrons, and getting scattered). Once the universe cooled down enough, electrons and protons formed hydrogen and helium atoms, which were transparent to photons. The photons could travel freely, and we see the ones that are reaching us just now (on a sphere). In fact, if you are able to tune an (old) television set between two channels, CMB forms 1% of the static. Because of the expansion of the universe, the radiation has been greatly red-shifted, and is now in the microwave region.

Its properties have been called “the discovery of the (twentieth) century, if not of all time”, by Stephen Hawking.

Among early applications of CMB was the confirmation of the Big Bang theory, based upon the observation that the spectrum of the radiation is that of a homogeneous black-body.

The temperature is nearly homogeneous, around 2.7 K, but with temperature fluctuations on the order of 10^{-4} degrees. These fluctuations, though small, form the basis for most of the analysis of the CMB. The regions of higher temperature corresponded roughly, at the time of CMB emission, to regions of higher density. The density fluctuations are of exactly the right amplitude to form the large-scale structures we see today. But physicists reason that there must be a lot of dark matter, or there would not have been time to form those structures.

From the analysis of CMB, physicists are able to estimate fundamental constants, such as the energy density (which is used to determine if the universe has spatial curvature; so far it appears to be very close to flat); the baryon density; the dark matter density; and the cosmological constant.

CMB also has a *polarization*, which has been much more difficult to measure, since it is much fainter than the temperature fluctuations. However, the Planck satellite, launched in May 2009, is presently measuring it very accurately, and the results will be made available in 2012. Physicists very much want to analyze

the polarization data, for two main reasons. First, one needs to understand that, a few hundred million years after the Big Bang, extremely massive and bright stars formed, which then reionized many of the hydrogen and helium atoms, which resulted in about 17% of the CMB photons being scattered anew by a thin fog of ionized gas. This resulted in a large part of the polarization of the CMB. The figures we have quoted here – “a few hundred million years” and “17 percent” are the estimates made based on the polarization data obtained by the earlier WMAP satellite, which measured the CMB polarization rather imprecisely. Planck should improve our understanding of the reionization period greatly, and will present important experimental data for theorists. Second, one expects evidence of gravitational waves in the polarization data. Gravitational waves have never been directly observed before.

Let us now briefly explain some of the mathematical ideas in the statistical analysis of CMB, and why spherical wavelets and spherical spin wavelets are considered a useful tool therein.

We view CMB as a *single* observation of a *second-order isotropic* random field.

Let’s just explain what this means for temperature; polarization is similar but a little more complicated. For a fuller discussion, see [3] and [11–13]. The temperature $T(x)$ at a point x on S^2 is thought of as a random quantity.

After subtracting off a constant, assume it has mean 0.

To say that the field is second-order isotropic means that covariances are invariant under rotation:

$$E(T(Rx)T(Ry)) = E(T(x)T(y))$$

for all $x, y \in S^2$, and all $R \in SO(3)$. Here E as usual denotes expectation, over an unknown probability space.

Let’s understand what second-order isotropy would mean if $\{T(x) : x \in S^2\}$ were in fact a *Gaussian family*.

Then all joint probability distributions are also Gaussian, and are completely determined by the covariances $E(T(x)T(y))$.

The second order isotropy condition

$$E(T(Rx)T(Ry)) = E(T(x)T(y))$$

(for all x, y, R) is then equivalent to the *isotropy* condition, which is that for any n , any $x_1, \dots, x_n \in S^2$, and any $R \in SO(3)$, the joint probability distribution of $T(x_1), \dots, T(x_n)$ is the same as that of $T(Rx_1), \dots, T(Rx_n)$.

Now let us expand $T(x)$ in spherical harmonics

$$T(x) = \sum_{l=0}^{\infty} \sum_{m=-l}^l A_{lm} Y_{lm}(x)$$

where now the A_{lm} are also random variables with mean 0. A crucial fact is that for each l ,

$$C_l := E(|A_{lm}|^2)$$

is *independent* of m . Moreover, the A_{lm} are all independent. This is a consequence of the second-order isotropy assumption

$$E(T(Rx)T(Ry)) = E(T(x)T(y))$$

and the irreducibility of the action of $SO(3)$ on the \mathcal{H}_l . (Examine the covariance matrices $\{E(A_{lm}\bar{A}_{lm'})\}_{m,m'}$, and apply Schur's lemma.)

The $\{C_l\}$ are called the *power spectrum*; once one knows them, one can verify or discount the physical predictions of various theories. Theorists long expected the graph of C_l versus l to have peaks at certain values of l . They have the following significance: the hot plasma that we discussed before had, in essence *sound waves*: here compression arose from gravitational attraction, while rarefaction arose from the pressure from the collisions. The first peak in the graph of the power spectrum arises from the *fundamental frequency* of these sound waves; succeeding peaks arise from overtones. The locations and sizes of the peaks are what give rise to the wealth of physical information we have alluded to.

The crucial question, then, is: how is one to obtain this graph of the power spectrum? All one knows is that $C_l := E(|A_{lm}|^2)$, where the expectation is taken over an unknown probability space. But the A_{lm} are all independent, so for l large one expects, by the *law of large numbers* to get a good estimate for C_l as

$$C_l \sim \frac{1}{2l+1} \sum_{m=-l}^l |A_{lm}|^2$$

where now we use the *observed* A_{lm} .

This process is similar to rolling a die many times, seeing if the average is about 3.5, and then concluding that the die is probably fair. This is the original procedure, used by physicists, for estimating the C_l .

So where do wavelets come in? The problem with using the observed spherical harmonic coefficients A_{lm} is that **we can't actually calculate them**. A large portion of the CMB data is contaminated (in an area called the “foreground” or “mask”), by competing emissions from the Milky Way and from other sources. The spherical harmonics, and the spin spherical harmonics, like the $e^{im\theta}$ on the circle, are supported all over the sphere, so we can't actually take the inner product $A_{lm} = \langle F, {}_sY_{lm} \rangle$.

One possibility is to estimate the foreground contamination arising from the Milky Way by other methods, and to subtract it from the observed data. However, especially for the faint polarization data, this is problematic. If the A_{lm} are corrupted, then they are no longer uncorrelated, so no law of large numbers applies.

Instead, one can take inner products with the well-localized spherical wavelets, and almost completely avoid the mask.

Recall that our wavelets arise from the kernel of $f(t^2\Delta_s)$. (Here, $s = 0$ for temperature and $s = 2$ for polarization.) The authors' contribution is to the study of polarization (Geller–Marinucci [13], where spin wavelets are first constructed, and the subsequent [11, 12, 17]), and to the use of Mexican needlets in the study of CMB temperature [15, 16, 26, 33]. Earlier work on CMB temperature, using band-limited f , was done by Baldi, Kerkycharian, Marinucci, and Picard [3], by using the spherical needlets of Narcowich, Petrushev, and Ward [28, 29].




Fig. 12.2 Reconstruction of polarization trial data, using spin wavelets

The question now arises as to which f to take. One wants several properties:

- The frames should be well-localized in frequency (essentially a linear combination of a small number of spherical harmonics ${}_sY_{lm}$ (with l restricted to a small interval). That way, we will be able to obtain, if not the C_l , then something quite similar to them.
- The frames should be well-localized in space, so we can ignore the missing data when we take inner products, as long as we stay away from the mask.
- The frame coefficients should at least be *asymptotically uncorrelated* in the high-frequency limit (large l), so that we can still apply the law of large numbers.

The first condition is automatic from our choice of f . The second condition follows from Theorem 12.2.

As for the third condition, say first $s = 0$. For band-limited f [3, 24, 28], and for $f(u) = ue^{-u}$ (the Mexican needlet situation) [15, 23, 26, 33] all three conditions are satisfied under mild assumptions. In fact [33], Mexican needlets appear to have numerical advantages for both the second and third requirements.

For general s , all three conditions are satisfied for band-limited f [13], under mild assumptions. For $s = 2$ (the polarization situation), one has (see Fig. 12.2) near-perfect reconstruction of Q and U for trial data, even at a very slight distance from the mask. Figure 2 is taken from [11], where it appears in color. As shown here in black and white, “blue”, the shade at the left of the greyscale denotes almost perfect reconstruction, while “red”, the shade at the right of the greyscale, denotes poor reconstruction. Red appears only very close to the mask.

References

1. J.-P. Antoine and P. Vandergheynst, *Wavelets on the sphere: a group-theoretic approach*, Applied and Computational Harmonic Analysis, **7** (1999), 262–291.
2. J.-P. Antoine and P. Vandergheynst, *Wavelets on the sphere and other conic sections*, Journal of Fourier Analysis and its Applications, **13** (2007), 369–386.

3. P. Baldi, G. Kerkycharian, D. Marinucci, D. Picard, *Asymptotics for Spherical Needlets*, Annals of Statistics, **37** (2009), 1150–1171.
4. I. Bogdanova, P. Vanderghelynst, J.-P. Antoine, L. Jacques and M. Morvidone, *Stereographic wavelet frames on the sphere*, Appl. Comput. Harmon. Anal. **19** (2005), 223–252.
5. M. Cruz, L. Cayon, E. Martinez-Gonzalez, P. Vielva, J. Jin, *The non-Gaussian cold spot in the 3-year WMAP data*, Astrophysical Journal **655** (2007), 11–20.
6. M. Cruz, L. Cayon, E. Martinez-Gonzalez, P. Vielva, P. *The non-Gaussian cold spot in WMAP: significance, morphology and foreground contribution*, Monthly Notices of the Royal Astronomical Society **369** (2006), 57–67.
7. I. Daubechies, *Ten Lectures on Wavelets*, Philadelphia, Pennsylvania, 1992.
8. M. Frazier and B. Jawerth, *Decomposition of Besov Spaces*, Ind. Univ. Math. J. **34** (1985), 777–799.
9. W. Freeden, T. Gervens and M. Schreiner, *Constructive Approximation on the Sphere, With Applications to Geomathematics*, Clarendon Press, Oxford (1998).
10. W. Freeden and M. Volker, *Multiscale Potential Theory*, Birkhauser, Boston (2004).
11. D. Geller, F.K. Hansen, D. Marinucci, G. Kerkycharian and D. Picard, *Spin needles for cosmic microwave background polarization data analysis*, Physical Review D, D78:123533 (2008), arXiv:0811.2881
12. D. Geller, X. Lan and D. Marinucci, *Spin needlets spectral estimation*, Electronic Journal of Statistics **3** (2009), 1497–1530, arXiv:0907.3369.
13. D. Geller and D. Marinucci, *Spin wavelets on the sphere*, Journal of Fourier Analysis and Applications, **16** (2010), 840–884, arXiv:0811.2835.
14. D. Geller and A. Mayeli, *Continuous wavelets and frames on stratified Lie groups I*, Journal of Fourier Analysis and Applications **12** (2006), 543–579.
15. D. Geller and A. Mayeli, *Continuous wavelets on compact manifolds*, Math. Z. **262** (2009), 895–927.
16. D. Geller and A. Mayeli, *Nearly tight frames and space-frequency analysis on compact manifolds*, Math. Z. **263** (2009), 235–264.
17. D. Geller and A. Mayeli, *Nearly tight frames of spin wavelets on the sphere*, to appear in Sampling Theory in Signal and Image Processing (2011), arXiv:0907.3164.
18. D. Geller and A. Mayeli, *Besov spaces and frames on compact manifolds*, Indiana University Math Journal **58** (2009), 2003–2042.
19. D. Geller and I. Pesenson, *Band-limited Localized Parseval frames and Besov spaces on Compact Homogeneous Manifolds*, to appear, Journal of Geometric Analysis, arXiv:1002.3841, DOI: 10.1007/s12220-010-9150-3.
20. Hansen, F.K., Cabella, P., Marinucci, D., Vittorio, N., *Asymmetries in the local curvature of the WMAP data*, Astrophysical Journal Letters (2004), p. L67–L70.
21. G. Kerkycharian, P. Petrushev, D. Picard, and Yuan Xu, *Decomposition of Triebel-Lizorkin and Besov spaces in the context of Laguerre expansions*, J. Funct. Anal. **256** (2009), 1137–1188.
22. G. Kyriazis, P. Petrushev, and Yuan Xu, *Decomposition of weighted Triebel-Lizorkin and Besov spaces on the ball*, Proc. London Math. Soc. **97** (2008), 477–513.
23. X. Lan, D. Marinucci, *On The Dependence Structure of Wavelet Coefficients for Spherical Random Fields*, Stochastic Processes and their Applications, **119** (2009), 3749–3766.
24. D. Marinucci, D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkycharian, P. Natoli, D. Picard, N. Vittorio, *Spherical Needlets for CMB Data Analysis*, Monthly Notices of the Royal Astronomical Society **383** (2008), 539–545, arXiv: 0707.0844.
25. A. Mayeli, *Discrete and continuous wavelet transformation on the Heisenberg group*, Ph.D thesis, Technische Universität München, 2006.
26. A. Mayeli, *Asymptotic Uncorrelation for Mexican Needlets*, Journal of Mathematical Analysis and Applications, **363**, 2010, 336–344.
27. J.D. McEwen, P. Vielva, Y. Wiaux, R.B. Barreiro, L. Cayon, M.P. Hobson, A.N. Lasenby, E. Martinez-Gonzalez, J. Sanz, *Cosmological applications of a wavelet analysis on the sphere*, Journal of Fourier Analysis and its Applications **13** (2007), 495–510.

28. F.J. Narcowich, P. Petrushev and J. Ward, *Localized tight frames on spheres*, SIAM J. Math. Anal. **38** (2006), 574–594.
29. F.J. Narcowich, P. Petrushev and J. Ward, *Decomposition of Besov and Triebel-Lizorkin spaces on the sphere*, J. Func. Anal. **238** (2006), 530–564.
30. E. T. Newman and R. Penrose, *Notes on the Bondi-Metzner-Sachs Group*, J. Math. Phys **7** (1966) 863–870.
31. G. Patanchon, J. Delabrouille, J.-F. Cardoso, P. Vielva, *CMB and foreground in WMAP first-year data*, Monthly Notices of the Royal Astronomical Society **364** (2005), 1185–1194.
32. P. Petrushev and Yuan Xu, *Decomposition of spaces of distributions induced by Hermite expansions*, J. Fourier Anal. and Appl. **14** (2008), 372–414.
33. S. Scodeller, O. Rudjord, F.K. Hansen, D. Marinucci, D. Geller, A. Mayeli, *Introducing Mexican needlets for CMB analysis: Issues for practical applications and comparison with standard needlets*, arxiv:1004.5576.
34. G.F. Smoot, G. F., C.L. Bennett, C. L., A. Kogut, E.L Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P.D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S.S. Meyer, S.H. Moseley, T. Murdock, L. Rokke, R.F. Silverberg, L. Tenorio, R. Weiss, D.T. Wilkinson, *Structure in the COBE differential microwave radiometer first-year maps*, Astrophysical Journal, Part 2 - Letters, **396**, no. 1, (1992) pp. L1–L5.
35. R. Strichartz, *A functional calculus for elliptic pseudodifferential operators*, Amer. J. Math **94** (1972), 711–722.
36. M. Taylor, *Pseudodifferential Operators*, Princeton University Press, (1981).
37. P. Vielva, E. Martínez-González, J.E. Gallegos, L. Toffolatti, J.L. Sanz, *Point source detection using the Spherical Mexican Hat Wavelet on simulated all-sky Planck maps*, Monthly Notice of the Royal Astronomical Society, **344**, Issue 1 (2003) 89–104.
38. Y. Wiaux, J.D. McEwen, P. Vielva, *Complex data processing: fast wavelet analysis on the sphere*, Journal of Fourier Analysis and its Applications, **13** (2007), 477–494.

Chapter 13

Wavelets, a Numerical Tool for Atmospheric Data Analysis

Parick Fischer and Ka-Kit Tung

Abstract Multiresolution methods, such as the wavelet decompositions, are increasingly used in physical applications where multiscale phenomena occur. We present in this review paper two applications illustrating that the classical 30-year-old continuous wavelets are still used and useful in physics.

In the first part, using a Continuous Wavelet Transform (CWT) for the determination of local QBO (Quasi-Biennial Oscillation) period, we reexamine the previous finding that the period of the QBO in the lower stratosphere is longer during solar minima.

In the second part, we use a wavelet-based multifractal approach to describe qualitatively and quantitatively the complex temporal patterns of atmospheric data. Time series of geopotential height are used in this study.

More detailed versions of these results have already been published in dedicated papers.

13.1 QBO-Period Modulation by the Solar Cycle

13.1.1 Presentation

The Quasi-Biennial Oscillation (QBO) is a dominant oscillation of the equatorial stratospheric zonal wind, whose period is irregular but averages to about 28 months. The classical mechanism of QBO attributes the period of the QBO to internal interactions between the waves and the mean flow in the equatorial stratosphere. Whether the QBO's period is affected by external forcing, such as the 11-year variation in the solar radiation (especially its variation in the UV component), is an intriguing open question.

P. Fischer (✉)

Institut de Mathématiques de Bordeaux, Université Bordeaux 1, 33405 Talence Cedex, France
e-mail: patrick.fischer@math.u-bordeaux1.fr

In this part, we shall reexamine the possibility of a decadal solar-cycle modulation of the period of the QBO using an objective method and the longest record available. We used in situ dataset of near equatorial wind at 50 hPa distributed by the Stratospheric Research Group at the Free University of Berlin (FUB), here updated by B. Naujokat to span from 1953 to 2007.

Calculating the period of the QBO has been a subjective procedure. It has usually involved visually determining when a descending westerly (or easterly) first crosses zero at a particular level and when it later goes back above zero. Such a procedure is sensitive to calibration errors and monthly averaging. The use of Continuous Wavelet Transform, which can determine the local intrinsic period of an oscillation, gives a more objective method that is not sensitive to the location of the zero-wind line.

13.1.2 Review on Continuous Wavelets

We recall here only the main ideas underlying the wavelet theory and we refer to [14] for a complete description. Any temporal signal, which can be seen as a one dimensional mathematical function, can be represented by a sum of fundamental functions called basis functions. The most famous example, the Fourier series,

$$s(t) = \sum_{k=-\infty}^{+\infty} c_k e^{ikt} \quad (13.1)$$

is valid for any 2π -periodic function sufficiently smooth. Each basis function, e^{ikt} is indexed by a parameter k which is related to a frequency. In (13.1), $s(t)$ is written as a superposition of harmonic modes with frequencies k . The coefficients c_n are given by the integral

$$c_k = \frac{1}{2\pi} \int_0^{2\pi} s(t) e^{-ikt} dt. \quad (13.2)$$

Each coefficient c_k can be viewed as the average harmonic content of $s(t)$ at frequency k . Thus the Fourier decomposition gives a frequency representation of any signal. The computation of c_k is called the decomposition of s and the series on the right hand side of (13.1) is called the reconstruction of s .

Although this decomposition leads to good results in many cases, some disadvantages are inherent to the method. One of them is the fact that all the information concerning the local time behavior of the signal is lost in the Fourier decomposition. For instance, a discontinuity or a localized high variation of the frequency will not be described by the Fourier representation in any intuitive or useful manner. The underlying reason lies in the nature of complex exponential functions used as basis functions. They are global functions that span the entire data record and differ only with respect to frequency.

Like the complex exponential functions of the Fourier decomposition, wavelets can be used as basis functions for the representation of a signal. But, unlike the complex exponential functions, they are able to restore the temporal information as well as the frequency information. The wavelet transform is a convolution product of a data sequence with the compressed (or dilated) and translated version of a basis function ψ called the mother wavelet. The scaling and translation are performed by two parameters: the scale parameter a dilates or compresses the mother wavelet to various resolutions and the translation parameter b moves the wavelet all along the sequence:

$$WT_s(b, a) = \frac{1}{\sqrt{a}} \int_{-\infty}^{+\infty} s(t) \psi^* \left(\frac{t-b}{a} \right) dt, \quad a \in \mathbb{R}^{+*}, b \in \mathbb{R}. \quad (13.3)$$

This definition of the wavelet transform leads to an invariant L^2 measure, and thus conserves the energy ($\|s\|_2 = \|WT_s\|_2$). A different normalization could be used leading to a different invariant.

A possibility is to construct the set $\{\psi_{a,b}(t)\}_{a \in \mathbb{R}^*, b \in \mathbb{R}}$ from a function $g(x)$ by translating and modulating it:

$$\psi_{a,b}(t) = g(t-b) e^{iat}, \quad (13.4)$$

where $g(t)$ is a window function. In spite of the improvement brought by this “pseudo-spectral” representation, this transformation is still not adapted to describing accurately functions which exhibit high local variations. To overcome this disadvantage (a fixed size window function), analyzing functions with time support widths adapted to their frequency are needed.

The idea is to apply dilations on top of translations previously introduced. Starting with a function ψ well localized in time and frequency spaces, a family of analyzing functions can be constructed:

$$\psi_{a,b}(t) = |a|^{-1/2} \psi \left(\frac{t-b}{a} \right). \quad (13.5)$$

where b is a time parameter and a is analogous to a period. The initial function ψ is called the mother wavelet and has to verify the following condition:

$$\int d\xi \frac{|\hat{\psi}|^2}{|\xi|} = K < \infty. \quad (13.6)$$

This condition means that any oscillating function localized in both spaces and whose integral over the whole space \mathbb{R} is null can be used as a mother wavelet. Similar to the definition of the inverse Fourier transform, it is also possible to define a reconstruction formula that allows one to rewrite $s(t)$ as an expansion on the corresponding basis. The coefficients $WT_s(b, a)$ defined in (13.3) give a time-scale representation of the initial signal allowing the detection of transient components or singularities.

The results presented in this part have been obtained using the tenth derivative of the Gaussian as a mother wavelet. A link between the Fourier period and the wavelet scale can be derived analytically for each mother wavelet. The results are given here in a time-period representation rather than the usual time-frequency representation of the wavelet theory.

13.1.3 Results and Discussion

13.1.3.1 QBO-period variation

Figure 13.1 shows the local period of the QBO oscillation as determined by applying the CWT to the FUB data at 50 hPa for the period 1953–2005. This radiosonde dataset of near equatorial zonal wind consists of observations at Canton Island (January 1953 – August 1967), Gan, Maldives (September 1967 – December 1975) and Singapore (since January 1976). It is the extended version of that used by [17] and [9]. It shows decadal variations around the mean period of 28 months. Because of this variation of periods, a Fourier analysis would give a broad spectrum of QBO period but is unable to locate the times with long or short periods. The color




Fig. 13.1 Local period in months of the 50 hPa FUB wind as determined by the CWT. The *darker region* denotes location of higher amplitude of the zonal wind. The *dashed line* traces the location of the maximum amplitude. Superimposed, in *solid line*, is the sunspot number (monthly averaged), which is a proxy for the solar cycle radiative flux variability

scheme in this figure shows the amplitude of the wavelet coefficients, with darker color indicating higher amplitudes. The maximum amplitude of these wavelet coefficients is marked in a dashed line in Fig. 13.1. This marks the dominant period of the equatorial zonal wind, i.e. the period associated with the most amplitude (or kinetic energy). This is the period we will be focusing on. Below it we also plot the sunspot number as a function of years, which is used as a proxy for the 11-year solar-cycle flux.

Figure 13.1 shows that, consistent with Salby and Callaghan, the period of the QBO reaches its maximum during the solar min of 1965, when the dominant period is 33 months, the solar min of 1976, when the dominant period is 30 months, and the solar min of 1986, when the dominant period is 31 months. Other than these three solar minima mentioned by Salby and Callaghan, however, the anti-correlation with the solar cycle breaks down. In the solar min of 1997, the dominant QBO period reaches a low of close to 25 months, consistent with the finding of Hamilton. Going forward in time, the correlation is the reverse of that of Salby and Callaghan. That is, during solar max, the QBO period is longer, while during solar min the QBO period is shorter. The in-phase relationship appears to commence around 1991. Prior to 1957, the period variation is also approximately in-phase with the solar cycle, as Hamilton already pointed out. Over the almost six cycles spanned by the FUB data, three cycles show anti-correlation of QBO period with solar flux, while the other two and a half cycles show in-phase correlation. As a consequence, the correlation coefficient between the two is close to zero (-0.05) for the long record of 1953–2005. It is intriguing to note the alternate correlation and anti-correlation of the QBO period with the solar cycle, which is a different behavior than two curves not related to each other at all that could also give a zero correlation coefficient. The calculation is repeated with the ERA-40 data. The result is very close to what we have shown here using the FUB data for the period of overlap.

13.1.3.2 Comparison with the subjective method

Figure 13.2 shows in the upper panel a comparison of the period as determined objectively using CWT and the subjectively determined period by measuring the length between successive zero crossing of the zonal wind. It turns out the result for the QBO period obtained by the subjective method is very different depending on whether one defines the full QBO period as easterly plus westerly, or as westerly plus easterly. Our method turns out to be consistent with the average of these two definitions, provided that monthly averages are used in the subjective method. The easterly plus westerly period appears to have a few more oscillations in 1980–1990. By comparing with the objectively determined period, one can perhaps attribute it to an artifact and not to volcanoes or the solar cycle.

On the lower panel of Fig. 13.2, the “period” of the westerly phase and that of the easterly phase are separately determined by the subjective method. These are consistent with the results of [17] and [9], but for the longer data record. It shows a decadal variation of the westerly period that track quite closely that of the full

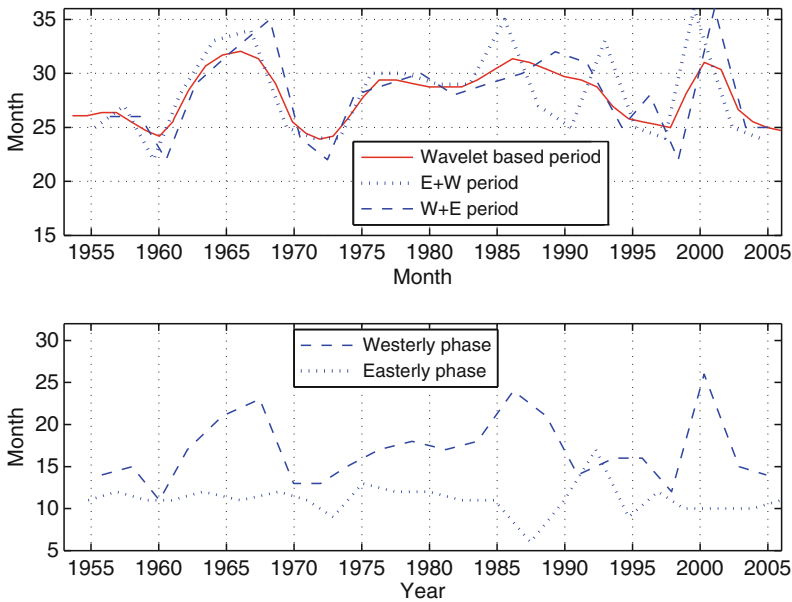


Fig. 13.2 *Upper panel:* A comparison of the QBO period as determined by the CWT method (in solid line) and that determined subjectively by adding the westerly phase period followed by the easterly phase period (in dashed line), and by adding the easterly phase period followed by the westerly phase period (in dotted line). *Lower panel:* The westerly phase duration (in dashed line) and the easterly phase duration (in dotted line)

QBO period as determined by the CWT. Therefore our conclusion that there is no correlation of the QBO period with the solar cycle also applies to the westerly phase separately. The correlation coefficient of the westerly-phase period variation with the sunspot time series for the full period of 1953–2005 is -0.10 .

Hamilton and Hsieh [10] proposed using the circular nonlinear principal component analysis to objectively analyze and characterize the quasi-periodic QBO oscillation. They found that a single time series of the QBO phase can be found for data at all levels. Although their method is not specifically aimed at studying the frequency variation of the QBO, a period variation similar to our CWT result was obtained in their Fig. 10, but with much high-frequency irregular oscillations. Hamilton and Hsieh [10] also concluded, based on their shorter record, that there is “no clear connection with the 11-year solar cycle”.

13.1.3.3 Behavior at different pressure levels

To answer the question of whether the QBO period changes with height, Fig. 13.3 shows the QBO period obtained the same way as in Fig. 13.1 using the FUB data for various pressure levels from 70 to 15 hPa. It shows the same prominent decadal




Fig. 13.3 The QBO period as determined by the CWT method using the FUB data for various pressure levels in the lower stratosphere

variation at all levels and that the difference are small with respect to height. Minor exceptions exist at the higher levels and at 70 hPa prior to 1958. These calculations were also repeated with the ERA-40 data, and the results were similar to those obtained using the FUB data for the period of overlap.

While the period variation of the QBO is almost the same at all heights in the lower stratosphere, above 30 hPa it is the easterly period variation that is responsible for most of the variation of the whole QBO period, but below 30 hPa it is the westerly period that controls the whole QBO period variation. Figure 13.4 shows that, interestingly, the easterly period at 15 hPa and the westerly period at 50 hPa vary synchronously. The amplitude of the variation is also about the same, from 12 to 23 months. This observational result can be understood as follows. When the mean equatorial upwelling is strong it slows down the descent of the easterlies. The effect being more noticeable on the easterlies than on the westerlies as explained in the Introduction. Above 30 hPa, there is no stalling of the easterlies. The slower descent of the easterlies then gives a longer easterly period. Below the level of stalling of the easterlies (near 30 hPa), however, there is no easterlies. Instead the westerlies at those levels persist without being replaced by the descending easterlies. This description explains the differing behavior of easterlies and westerlies above and below the stalling level, while the whole period of the QBO remains the same at these levels.




Fig. 13.4 The period of the westerly phase at 50 hPa (in *solid line*) and that of the easterly phase at 15 hPa as a function of year

13.1.4 Conclusion

The longer equatorial zonal wind dataset from Free University of Berlin (1953–2005) spans almost six solar cycles. We have found that during three of the cycles the period of the QBO is anti-correlated with the solar cycle, while in the remaining almost three cycles, there is correlation with the solar-cycle flux. Consequently, over the five and a half solar cycles the correlation coefficient is zero. The period previously considered by [17] contains three anti-correlated periods with one “straddling” period in 1992, which could be discounted as due to Pinatubo. With our longer record extending into 2005, when there has not been a major volcanic eruption since 1991, it becomes more difficult to attribute the “anomalous” behavior to volcanic aerosols. Our result strengthens that of [9] by using a longer data record, a more objective method of determining local period, and by showing that our conclusion is the same for all levels in the lower stratosphere. Our result however does not rule out the effect of solar cycle on the QBO period. It is rather intriguing to find that the variation of the QBO period is not random, but follows closely the variation of the solar cycle. It is only that the correlation completely reverses itself after three solar cycles. The possibility exists that it is the period of 1960s, 1970s, 1980s, and 1990s that is anomalous, with three major volcanic eruptions (Agung in 1963, El Chichon in 1982 and Pinatubo in 1991), and instead the QBO period

should actually be shorter at solar minimum and longer at solar maximum in the absence of volcanic aerosols. Angel [1] argued that the prolongation of the westerly phase after the Agung eruption in 1963 was probably due to the aerosol heating; the temperature increase was also seen at 50 and 30 hPa at Balboa station. He also suggested that the higher altitude reached by the El Chichon aerosol in 1982 produced a “shielding” effect at 50 hPa, which might have prevented the aerosol heating at that altitude in 1982. The mechanism of solar cycle influence of the equatorial QBO remains yet to be discovered.

13.2 Wavelet Based Multifractal Formalism

13.2.1 *Presentation*

A wavelet transform can focus on localized signal structures with a zooming procedure that progressively reduces the scale parameter. The local signal regularity can be described by the decay of the wavelet transform amplitude across scales. Singularities can be detected by following the wavelet transform local maxima at fine scales. Finding the distribution of singularities in a multifractal signal is necessary for analyzing its properties. The so-called spectrum of singularity $D(h)$ measures the repartition of singularities having different Hölder regularity. The singularity spectrum $D(h)$ gives the proportion of Hölder h type singularities that appear in the signal. A fractal signal has only one type of singularity, and its singularity spectrum is reduced to one point.

We have applied the wavelet-based multifractal approach to the analysis of two sets of atmospheric data. The first set consists of the monthly averages of the NCEP Daily Global Analyzes data [16]. They correspond to time series of geopotential height from January 1948 to June 2005. A spatial average from 60 to 90° N is performed at 17 levels, from 10 down to 1,000 hPa. Then the annual cycle is removed by subtracting for each month the corresponding mean in order to focus our study onto the anomalies. In such way, we will be able to detect and to describe the singularities present in the signal. The second set of data consists of the Northern Annular Modes (NAM) at 17 levels from the stratosphere down to the surface level from January 1958 to July 2006 provided by Baldwin [6]. The results obtained with this second set of data are not given in this paper, and the reader interested in atmospheric sciences can find them in [7, 8].

13.2.2 *Continuous Wavelets and the Multifractal Formalism*

The wavelet-based multifractal formalism has been introduced in the nineties by Mallat [13, 14], Arneodo [2–4], Bacry [5] and Muzy [15]. A wavelet transform can focus on localized signal structures with a zooming procedure that progressively

reduces the scale parameter. Singularities and irregular structures often correspond to essential information in a signal. The local signal regularity can be described by the decay of the wavelet transform amplitude across scales. Singularities can be detected by following the wavelet transform local maxima at fine scales.

The strength of a singularity of a function is usually defined by an exponent called Hölder exponent. The Hölder exponent $h(t_0)$ of a function s at the point t_0 is defined as the largest exponent such that there exists a polynomial $P_n(t)$ of order n satisfying:

$$|s(t) - P_n(t - t_0)| \leq C|t - t_0|^{h(t_0)}, \quad (13.7)$$

for t in a neighborhood of t_0 . The order n of the polynomial P_n has to be as large as possible in (13.7). The polynomial P_n can be the Taylor expansion of s around t_0 . If $n < h(t_0) < n + 1$ then s is C^n but not C^{n+1} . The exponent h evaluates the regularity of s at the point t_0 . The higher the exponent h , the more regular the function s . It can be interpreted as a local measure of ‘burstiness’ in the time-series at time t_0 . A wavelet transform can estimate this exponent by ignoring the polynomial P_n . A transient structure or ‘burst’ is generally wavelet-transformed to a superposition of wavelets with the same centre of mass and wide range of frequencies.

In order to evaluate the Hölder exponent, we have to choose a mother wavelet with $m > h$ vanishing moments:

$$\int_{-\infty}^{\infty} t^k \psi(t) dt, \quad (13.8)$$

for $0 \leq k < m$. A wavelet with m vanishing moments is orthogonal to polynomials of degree $m - 1$. Since $h < m$, the polynomial P_n has a degree n at most equal to $m - 1$ and we can then show that:

$$\int_{-\infty}^{+\infty} P_n(t - t_0) \psi^* \left(\frac{t - b}{a} \right) dt = 0. \quad (13.9)$$

Let us assume that the function s can be written as a Taylor expansion around t_0 :

$$s(t) = P_n(t - t_0) + C|t - t_0|^{h(t_0)} \quad (13.10)$$

We then obtain for its wavelet transform at t_0 :

$$WT_s(t_0, a) = \frac{1}{\sqrt{a}} \int_{-\infty}^{+\infty} C|t - t_0|^{h(t_0)} \psi^* \left(\frac{t - t_0}{a} \right) dt \quad (13.11)$$

$$= C|a|^{h(t_0)+\frac{1}{2}} \int_{-\infty}^{+\infty} |t'|^{h(t_0)} \psi(t') dt'. \quad (13.12)$$

We have the following power law proportionality for the wavelet transform of the singularity of $s(t_0)$:

$$|WT_s(t_0, a)| \sim a^{h(t_0)+\frac{1}{2}} \quad (13.13)$$

Then, we can evaluate the exponent $h(t_0)$ from a log-log plot of the wavelet transform amplitude versus the scale a .

However, we cannot compute the regularity of a multifractal signal because its singularities are not isolated. But we can still obtain the singularity spectrum of multifractals from the wavelet transform local maxima.

These maxima are located along curves in the plane (b, a) . This method, introduced by Arneodo et al. [3], requires the computation of a global partition function $Z(q, a)$. Let $\{b_i(a)\}_{i \in \mathbb{Z}}$ be the position of all maxima of $|WT_s(b, a)|$ at a fixed scale a . The partition function $Z(q, a)$ is then defined by:

$$Z(q, a) = \sum_i |WT_s(b_i, a)|^q. \quad (13.14)$$

We can then assess the asymptotic decay $\tau(q)$ of $Z(q, a)$ at fine scales a for each $q \in \mathbb{R}$:

$$\tau(q) = \liminf_{a \rightarrow 0} \frac{\log Z(q, a)}{\log a}. \quad (13.15)$$

This last expression can be rewritten as a power law for the partition function $Z(q, a)$:

$$Z(q, a) \sim a^{\tau(q)}. \quad (13.16)$$

If the exponents $\tau(q)$ define a straight line, then the signal is a monofractal, otherwise the signal is called multifractal: the regularity properties of the signal are inhomogeneous, and change with location.

Finding the distribution of singularities in a multifractal signal is necessary for analyzing its properties. The so-called spectrum of singularity $D(h)$ measures the repartition of singularities having different Hölder regularity. The singularity spectrum $D(h)$ gives the proportion of Hölder h type singularities that appear in the signal. A fractal signal has only one type of singularity, and its singularity spectrum is reduced to one point. The singularity spectrum $D(h)$ for any multifractal signal can be obtained from the Legendre transform of the scaling exponent $\tau(q)$ previously defined :

$$D(h) = \min_{q \in \mathbb{R}} \left(q \left(h + \frac{1}{2} \right) - \tau(q) \right). \quad (13.17)$$

Let us notice that this formula is only valid for functions with a convex singularity spectrum [14]. In general, the Legendre transform gives only an upper bound of $D(h)$ [11, 12]. For a convex singularity spectrum $D(h)$, its maximum

$$D(h_0) = \max_h D(h) = -\tau(0) \quad (13.18)$$

is the fractal dimension of the Hölder exponent h_0 .

Remark: When the maximum value of the modulus of the wavelet transform is very small, the formulation of the partition function given in (13.14) can diverge

for $q < 0$. A way to avoid this problem consists in replacing the value of the modulus of the wavelet transform at each maximum by the supremum value along the corresponding maxima line at scales smaller than a :

$$Z(q, a) = \sum_{l \in \mathcal{L}(a)} \left(\sup_{(t, a') \in l, a' < a} |WT_s(t, a)| \right)^q, \quad (13.19)$$

where $\mathcal{L}(a)$ is the set of all maxima lines l satisfying: $l \in \mathcal{L}(a)$, if $\forall a' \leq a, \exists (x, a') \in l$. The properties of this modified partition function are well described in [3].

13.2.3 Multifractal Analysis of Atmospheric Time Series

Depending on the application, there are various ways of computing the wavelet transform. For the purpose of compression for instance, an orthogonal wavelet transform on dyadic scales are generally used. For the study of fractals like in this present study, continuous wavelet transforms have been found to be efficient [5]. The mother wavelet has also to be chosen according to the application. When the time series do not have any characteristic scales, or when the goal is to identify discontinuities or singularities, a real mother wavelet has to be chosen. In this work, we use the N successive derivatives of a Gaussian function:

$$\psi(x) = \frac{d^N}{dx^N} e^{-x^2/2} \quad (13.20)$$

These functions are well localized in both space and frequency, and have N vanishing moments, as required for a multifractal analysis. The computations have been performed with $N = 1, 2, 4, 6, 8, 10$ but only the results for $N = 2$ are discussed in detail in this paper. The results for other values of N are very similar in denoting the absence of any polynomial component. Furthermore, the case $N = 2$ is generally used for fractal analysis and corresponds to the so-called Mexican Hat function.

13.2.3.1 Data setup

We have applied the wavelet-based multifractal approach to the analysis of two sets of atmospheric data. The first set consists of the monthly averages of the NCEP Daily Global Analyzes data [16]. They correspond to time series of geopotential height from January 1948 to June 2005. A spatial average from 60 to 90° N is performed at 17 levels, from 10 down to 1,000 hPa. Then the annual cycle is removed by subtracting for each month the corresponding mean in order to focus our study onto the anomalies. In such way, we will be able to detect and to describe the singularities present in the signal. Typical stratospheric and tropospheric representations are shown at 100 and 700 hPa in Figs. 13.5 and 13.6.




Fig. 13.5 100 hPa monthly anomalies (NCEP) from 60 to 90° N




Fig. 13.6 700 hPa monthly anomalies (NCEP) from 60 to 90° N

The second set of data consists of the Northern Annular Modes (NAM) at 17 levels from the stratosphere down to the surface level from January 1958 to July 2006 provided by Baldwin [6]. At each pressure altitude, the annular mode is the first Empirical Orthogonal Function (EOF) of 90-day low-pass filtered geopotential anomalies north of 20° N. Daily values of the annular mode are calculated for each pressure

altitude by projecting daily geopotential anomalies onto the leading EOF patterns. In the stratosphere annular mode values are a measure of the strength of the polar vortex, while the near-surface annular mode is called the Arctic Oscillation (AO), which is recognized as the North Atlantic Oscillation (NAO) over the Atlantic sector.

13.2.3.2 Numerical results

The wavelet decompositions obtained with the Mexican Hat function (second derivative of the Gaussian function) are given in Figs. 13.7 and 13.8. The wavelet transform consists in the calculation of a resemblance index between the signal and the mother wavelet (here the Mexican Hat function). If the signal is similar to itself at different scales, then the wavelet coefficients representation will be also similar at different scales. It can be easily noticed in Figs. 13.7 and 13.8 that the self-similarity generates a characteristic pattern. This representation is a good demonstration of how well the wavelet transform can reveal the fractal pattern of the atmospheric data. Based only on these representations, we cannot make any significant difference between the stratospheric and the tropospheric signals. But we will see in the sequel by studying the maxima lines of the wavelet transform that these two signals have a different singularity spectrum $D(h)$.

Based on the technical reasons presented in the previous section, the partition function is computed with the formulation given in (13.19) for q between -20 and 20 with a step size of 0.5 .

The first step in the computation of the partition function consists in the detection of the maxima lines of the modulus of the wavelet transform. The representation of




Fig. 13.7 Modulus of the wavelet transform of the 100 hPa signal




Fig. 13.8 Modulus of the wavelet transform of the 700 hPa signal




Fig. 13.9 Maxima lines of the modulus of the wavelet transform of the 100-hPa signal

these maxima lines, often called the “skeleton” of the wavelet transform, is given in Fig. 13.9 for the stratospheric signal. For the computation of the partition functions, only the maxima lines of length longer than 1 octave are kept in the summation in order to keep only the significant singularities. The two partition functions are given in Figs. 13.10 and 13.11. The steps that can be observed for negative values of q




Fig. 13.10 Partition function at 100 hPa




Fig. 13.11 Partition function at 700 hPa




Fig. 13.12 Singularity spectra of the 100 and 700 hPa signals

are due to the use of the supremum (otherwise, the computation of $Z(q, a)$ would diverge for negative q). We can remark that the slopes for negative q are different for the stratosphere and for the troposphere. Based on this simple remark, we can already predict that the shapes of the corresponding singularity spectra will be also different. We can expect a steeper down slope in the case of the troposphere.

The corresponding singularity spectra are given in Fig. 13.12. The large supports of the spectra prove that the signals are multifractal. A quasi-monofractal signal spectrum would lie on very few values, and a real monofractal signal spectrum would reduce to only one point.

As expected, the down slope corresponding to the negative values of q is steeper for the troposphere than for the stratosphere. The maximum of the spectra is obtained around $h = -0.29$ for the stratosphere and between $h = -0.22$ and $h = -0.23$ for the troposphere. We recall that the smaller the value h is, the stronger the singularities in the signal are.

So according to this first study, we can conclude that the singularities in the tropospheric signal are more singular than the singularities in the stratospheric signal. We can verify this first conclusion by computing the value of h where the maximum of $D(h)$ is obtained for the 17 levels from 10 down to 1,000 hPa. The results are given in Fig. 13.13. We can clearly detect two areas: the first one with h around -0.23 corresponds to the stratosphere and the second one with h around -0.29 corresponds to the troposphere. These results can be compared to the values h obtained for artificial uncorrelated data. We perform the same computations on signals of random numbers whose elements are uniformly distributed in the interval $(0, 1)$. The value of h found for random signals are around -0.4 . So with $h \sim -0.3$ or $h \sim -0.2$, the signals corresponding to atmospheric data are close to artificial uncorrelated data at these ranges of time periods.




Fig. 13.13 Evolution of h in function of the level




Fig. 13.14 Singularity spectra for few levels in the stratosphere and in the troposphere

The whole singularity spectra can also give some information to discriminate stratospheric data from tropospheric data. We can show that their supports are also different as can be noticed from Fig. 13.14. The stratospheric signals present broader spectra than the tropospheric signals indicating the presence of singularities over a larger spectrum.

The analysis performed on the monthly averages NCEP Data cannot give any information for periods smaller than a month. In order to get details on finer time periods, we performed the same kind of analysis on the daily NAM index. The corresponding results are given in [7, 8].

13.2.3.3 Discussion

In this part, we have discussed some issues relating to the estimation of the multifractal nature of atmospheric data using a wavelet-based method. Our study reveals the clear fractal pattern of the analyzed series and their different scaling characteristics. The results obtained with daily data (not shown here) show, in the case of the stratosphere, a short-range correlation behavior that occurs for short range of time scales. In the troposphere and in the same ranges of time, we found a much weaker correlation. The connection of the multiscaling properties of atmospheric data to the underlying physical dynamics falls beyond the scope of the present review paper. However, by using a two dimensional wavelet transform, we would like to extend our research from time series to spatial patterns of atmosphere analysis.

13.3 Conclusion

We have reviewed in this paper a few results obtained in two different problems related to atmospheric data analysis. Wavelets were developed independently in the fields of mathematics, quantum physics, geology, and electrical engineering. They are perfect numerical tools in analyzing physical situations where the signal contains discontinuities or sharp spikes, and they are especially adapted for studying multiscale phenomena in many physical applications. Twenty years after their formal introduction in 1989, the good old wavelet decompositions are still largely used and appreciated by scientific communities.

Acknowledgements The research was supported by the National Science Foundation, Climate Dynamics Program, under grant ATM-0332364, and the DGA (French Defense Department) under contract 06.60.018.00.470.75.01.

References

1. Angel J.K.: On the variation in period and amplitude of the quasi-biennial oscillation in the equatorial stratosphere 1951–1985, *Monthly Weather Rev.*, **114**, 2272–2278 (1986).
2. Arneodo, A., Grasseau, G., Holschneider, M.: Wavelet transform of multifractals, *Phys. Rev. Lett.*, **61**, 2281–2284 (1988).
3. Arneodo, A., Bacry, E., Muzy, J.F.: The thermodynamics of fractals revisited with wavelets, *Physica A*, **213**, 232–275 (1995).

4. Arneodo, A., Argoul, F., Bacry, E., Elezgaray, J., Muzy, J.F.: *Ondelettes, multifractales et turbulence*, Diderot Editeur, Paris, France (1995).
5. Bacry, E., Muzy, J.F., Arneodo, A.: Singularity spectrum of fractals signal from wavelet analysis: Exact results, *J. Stat. Phys.*, **70**, 635–674 (1993).
6. Baldwin, M.P., <http://www.nwra.com/resumes/baldwin/nam.php>
7. Fischer P, Tung K.K.: Wavelets, a numerical tool for multiscale phenomena: from two-dimensional turbulence to atmospheric data analysis., *J. Num. Anal. And Mod.*, **5**, 64–84 (2008).
8. Fischer P, Tung K.K., Wavelet-based Multifractal Analysis of atmospheric data, Draft version.
9. Hamilton, K.: On the quasi-decadal modulation of the stratospheric QBO period, *J. Climate*, **15**, 2562–2565 (2002).
10. Hamilton, K., Hsieh, W.: Representation of the quasi-biennial oscillation in the tropical stratospheric wind by nonlinear principal component analysis, *J. Geophys. Res.*, **10 D15**, 4232, 10.1029/2001JD001250 (2002).
11. Jaffard, S.: Multifractal formalism for functions Part I: Results valid for all functions, *SIAM J. Math. Anal.*, **28**, 944–970 (1997).
12. Jaffard, S.: Multifractal formalism for functions Part II: Self-similar functions, *SIAM J. Math. Anal.*, **28**, 971–998 (1997).
13. Mallat, S., Zhong, S.: Wavelet transform maxima and multiscale edges, in: R.M. B. et al. (Eds.), *Wavelets and their Applications*, Jones and Bartlett, Boston (1991).
14. Mallat, S.: *A wavelet tour of signal processing*, Academic Press, New York (1998).
15. Muzy, J.F., Bacry, E., Arneodo, A.: Wavelets and multifractal formalism for singular signals: application to turbulence data, *Phys. Rev. Lett.*, **67**, 3515–3518 (1991).
16. NOAA-CIRES Climate Diagnostics Center in Boulder, Colorado, USA, <http://www.cdc.noaa.gov>
17. Salby, M., Callaghan, P.: Connection between the solar cycle and the QBO: The missing link, *J. Climate*, **13**, 2652–2662 (2000).

Chapter 14

Denoising Speech Signals for Digital Hearing Aids: A Wavelet Based Approach

Nathaniel Whitmal, Janet Rutledge, and Jonathan Cohen

Abstract This study describes research developing a wavelet based, single microphone noise reduction algorithm for use in digital hearing aids. The approach reduces noise by expanding the observed speech in a series of implicitly filtered, shift-invariant wavelet packet basis vectors. The implicit filtering operation allows the method to reduce correlated noise while retaining low-level high-frequency spectral components that are necessary for intelligible speech. Recordings of speech in automobile road noise at signal-to-noise ratios of 0, 5, 10, 15, and 20 dB were used to evaluate the new method. Objective measurements indicate that the new method provides better noise reduction and lower signal distortion than previous wavelet-based methods, and produces output free from audible artifacts of conventional FFT methods. However, trials of the Revised Speech Perception in Noise test with the new algorithm showed no significant improvement in speech perception. Subsequent analysis has shown that the algorithm imposes physical attenuation on low-intensity components that mimics the perceptual effects of mild hearing loss.

14.1 Introduction

Unwanted acoustic noise is present in a variety of listening environments. Common examples include the background of conversational babble found in a cocktail party or a crowded restaurant; the broad spectrum noise produced by a loud air-conditioner or a jet engine at an airport, and the road noise heard in a car during highway driving. These noises can impair speech communication by “masking” (reducing the audibility of) nearby speech signals. Masking is particularly troublesome for hearing-impaired listeners, who have greater difficulty understanding speech in noise than normal-hearing listeners. Our research applied the mathematical theory of wavelet denoising as developed by Coifman and Saito [15, 60] and Donoho and

J. Cohen (✉)

Department of Mathematics, DePaul University, Chicago, Illinois, USA

e-mail: jcohen@math.depaul.edu

Johnstone [18, 19], to reducing unwanted noise for users of digital hearing aids. Specifically, we evaluated the utility of wavelets and wavelet packets in algorithms designed to remove the noise from noisy speech signals.

The final version of our algorithm was the result of a series of successive experiments in which the algorithm was adjusted to overcome a number of problems, including; the selection of a basis that best discriminated between noise and speech; choosing an optimal number of coefficients from which to reconstruct the denoised signal; removing the artificial sounds introduced by the processing; creating an algorithm robust enough to be able to respond to different types of noise.

The main body of the paper begins with a description of the original experiments that were performed on a set of nine noisy speech files. The next section is a detailed description of the algorithm we used to remove background noise. The following two sections are performance evaluations of the algorithm. Section 14.3 is devoted to objective measures of the performance of the algorithm. Using spectrograms and signal-to-noise ratios (SNR), we show improvements in our algorithm as compared to other forms of processing. Section 14.4 reports the results of testing the algorithm on a small group of hearing impaired subjects. These results produced intelligibility scores comparable to the algorithms with which it was compared. Section 14.5 uses the results of subsequent experiments to show that the hard thresholding used in our algorithm produces distortions similar to those produced by algorithms that simulate recruitment of loudness for non-impaired listeners. This partially explains how the algorithm produced improvements in objective measures without a corresponding improvement in intelligibility. In the final section, we discuss some of the lessons learned from our research.

Before proceeding to the description of our work, it is useful to review some of the other algorithms that have been tried to improve intelligibility. We also point out some of the difficulties that arise that are specific to the denoising of speech.

Poor intelligibility in background noise has been a long-standing problem for hearing aid users. Until about 25 years ago, intelligibility problems were related to the power and fidelity limitations of hearing aids themselves. A review of these constraints is given by Killion [39]. Subsequent improvements in the size, power consumption, and capability of digital signal processors led to the adoption of digital signal processing (DSP) technology in hearing aids. The first commercial digitally programmable hearing aids were introduced in the late 1980s, and followed in 1996 by two all-digital hearing aids: the Widex Senso and the Oticon Digifocus (Levitt [43]). Both digital aids had an immediate effect on the hearing aid market. A survey taken in 1997 revealed that programmable/digital aids already accounted for 13% of the market; this share increased to 28% by 2000 and to 51% by 2005 (Kochkin [42]). Nevertheless, surveys taken over the period of our research (Kochkin [40, 41]) indicated that only 25%–30% of hearing aid users remained satisfied with the performance of their hearing aids in noisy environments.

In anticipation of all-digital hearing aids, several 1980s-era research groups attempted to use DSP algorithms to enhance noisy speech for hearing aid users. These algorithms may generally be divided into two groups: multiple-microphone approaches and single microphone approaches.

Multiple-microphone approaches (Brey et al. [8]; Chabries et al. [10]; Schwander and Levitt [62]; Peterson et al. [54]; Greenberg and Zurek [33]), exploited the correlation between noise signals from spatially separated microphones to enhance noisy speech. While this approach was shown to improve intelligibility, the microphone spacings used in those studies rendered their algorithms impractical for commercial hearing aids.

The single microphone approaches of (Lim and Oppenheim [47]; Boll [4]; McAulay and Malpass [51]; Porter and Boll [55]; Ephraim and Malah [25]; Graupe et al. [30]; Ephraim et al. [26]; Kates [38]), relied on simple probabilistic models of speech and noise. These approaches were considered appropriate for digital hearing aids and other portable communication systems (e.g., cellular phones) for which spatially separated inputs were impractical. These noise reduction algorithms generally attempted to estimate temporal or spectral features of the original signal from observations of the noise-corrupted signal. These approaches met with mixed success, particularly at low SNRs, where the algorithms attenuated or distorted consonants which contribute to intelligibility. Numerous evaluations using both objective (Lim [46]; Makhoul and McAulay [49]; Boll [5]) and subjective (Tyler and Kuk [67]; Sammeth and Ochs [61]; Levitt et al. [45]) criteria indicated that these single-microphone algorithms failed to produce consistent improvements in intelligibility.

A promising alternative class of single-microphone noise reduction methods used the subspace approach, which projects segments of the noisy speech signal onto orthogonal “speech” and “noise” subspaces and discards the “noise” component. The speech subspace is constructed from high-energy vectors in the segment’s principal-component (or Karhunen–Loève) basis. One type of subspace algorithm was evaluated in a series of papers by Ephraim and Van Trees [27] and [28] and Ephraim et al. [29]. By using long segments of speech, Ephraim and Van Trees were able to use a discrete Fourier transform (DFT) to approximate principal-component bases in a computationally efficient manner. Their objective test results [28] showed that their approach eliminates the “musical noise” artifacts produced by earlier methods like spectral subtraction. Unfortunately, their subjective test results [29] indicate that (like earlier approaches), the subspace approach improves perceived speech quality without improving speech intelligibility.

It is possible that their algorithm’s performance may be affected by the poor time localization of DFT basis functions, which can limit the function’s usefulness in low order models of transient events (e.g. plosive consonants). Another factor may be the spectral and temporal properties of the 44 English phonemes (i.e. speech sounds). Vowels have a predictable quasi-periodic spectrum and may compress easily in a Fourier basis. Fricatives on the other hand, such as /s/, /sh/ or /f/, do not localize as well in frequency and are almost indistinguishable from noise; therefore, the algorithm tends to delete them along with the background noise. The balance of phonemes, that is the various families of consonants, are somewhere in between and are only partially removed by denoising. Vowels are higher in amplitude than consonants, so background noise tends to mask consonants to a greater extent. This situation is compounded by the further auditory degradation resulting from hearing loss.

The orthogonal basis functions produced by the discrete wavelet transform (Mallat [50]) overcome some limitations of the DFT basis, providing good spectral localization at low frequencies and good temporal resolution at high frequencies. Projection of a signal onto these basis functions may be accomplished efficiently by passing the signal through a tree-structured conjugate quadrature filter bank (Smith and Barnwell [65]). More flexible time–frequency localization may be obtained with the wavelet-packet transform (WPT) of Coifman and Wickerhauser [16], which allows the signal to be decomposed into subbands providing appropriate temporal and spectral resolution for a given application. Both the wavelet and wavelet-packet transforms demonstrate compression capabilities that rival those of principal component bases for many types of real world signals (Wornell [76]; Donoho [18]; Sinha and Tewfik [64]). This combination of properties motivated several researchers to implement subspace denoising with wavelet transforms. Evaluation of subspace denoising using shift-variant and shift-invariant versions of the discrete wavelet (Donoho and Johnstone [19]; Donoho et al. [20]; Coifman and Donoho [12]; Lang et al. [44]) and wavelet packet transforms (Coifman and Majid [14]; Saito [60]; Berger et al. [2]; Pesquet et al. [53]) demonstrated improvements in signal-to-noise ratio (SNR) ranging between 5 and 10 dB for signals with initial SNRs between 0 and 20 dB. Shift invariant algorithms (which implicitly average together several enhanced versions of the signal) provided additional improvements of as much as 7 dB at the cost of some computational efficiency. Of the methods described above, only one (Berger et al. [2]) was evaluated with speech; for this method, no objective or subjective data were reported.

The objective of our study was to evaluate wavelet-based subspace noise reduction for use in digital hearing aid applications. Results of a preliminary study are described; these are followed by specification and development of a new algorithm evaluated with both objective measures and listening results from trials of the Revised Speech Perception in Noise (R-SPIN) test (Bilger et al. [3]) with both normal acuity and impaired acuity listeners.

14.1.1 The First Algorithm

The first experiments were conducted with a software package that used an entropy criterion to find the coefficients for a wavelet packet best basis. The entropy criterion was chosen to pick the basis that concentrated the largest amount of energy in the fewest number of coefficients. With the coefficients placed in decreasing order by size, the best basis was the one that began with the steepest negative slope. Before applying the software to the signal, the experimenter was prompted for a percentage of coefficients that were to be retained in the reconstruction of the signal. The selection process then started with the largest coefficient and proceeded in descending order until the prescribed percentage of coefficients was reached.

Letting S denote the original signal, we let C_1 , the “coherent” signal, denote the signal reconstructed from the selected largest coefficients. The “noise” signal N_1

from the remaining coefficients associated with this selection of basis elements is the residual; that is,

$$S = C_1 + N_1.$$

The algorithm was then applied to the residual N_1 which resulted in the decomposition of N_1 into coherent and noise parts

$$N_1 = C_2 + N_2$$

Repeating the process yielded a sequence of C s and N s that was terminated as soon as the noise part was essentially free of any speech signal. The speech signal is then represented as the sum

$$C_1 + \cdots + C_n$$

of the coherent parts, all of which have presumably been separated from the noise.

The idea is that the speech is coherent and is concentrated in a few coefficients and the noise spreads out evenly among many coefficients. By repeating the procedure, the primary and secondary coherent parts of the signal can be peeled off, leaving the sequence of N_i s containing decreasing amounts of the speech part of the original signal. With this version of the software, the experimenter could adjust the percentage of the coefficients to keep at each stage of the algorithm.

14.1.2 The Original 9 Noisy Sentences

The original experiments were conducted at a computer lab in the Northwestern University Department of Electrical Engineering and Computer Science and the Department of Communication Sciences and Disorders. There were attempts to improve the intelligibility of 9 sentences chosen from the Harvard/IEEE sentence corpus [35], that had been recorded and corrupted by the addition of white noise.

The 9 phonetically balanced sentences were designed to minimize the ability of the listener to figure out the words from the context of the sentences. The noisy speech files were created by adding a file of randomly generated samples from a normal distribution to the original speech file. The signal-to-noise ratios of the sentences were sufficiently low that it was difficult for listeners with normal hearing to recognize any of the words, particularly on the first time hearing them. On the other hand, the noise level was sufficiently low that the words were readily recognizable to someone who knew the sentences beforehand. The nine sentences were:

1. The pipe began to rust while new.
2. Thieves who rob friends deserve jail.
3. Add the sum to the product of these three.
4. Open the crate but don't break the glass.
5. Oak is strong and also gives shade.
6. Cats and dogs each hate the other.

7. That hose can wash her feet.
8. Go and sit on the bed.
9. I read the news today.

As an indication of how well disguised the sentences were, out of a group of several listeners to the first noisy sentence, only one guessed that the second word was pipe. Not surprisingly, the one who correctly identified pipe was the only native English speaker.

The nine sentences were originally unknown to the person conducting the experiments. Of the nine sentences, it was only for the last one that the experimenter was able to identify the exact sentence after repeated applications of the algorithm. And even in this case it is difficult to say whether this success was due to the improved intelligibility or simply to the repeated hearing of the noisy sentence.

A more sophisticated version of the software was obtained from the group at Yale that enabled the user to specify the number of levels (iterations of the signal decomposition into coherent and noise parts) to run the algorithm and also gave some choice of wavelet to use. This provided additional flexibility and efficiency. But it also gave up some flexibility as it required the user to specify the same percent of coefficients to keep at each level of the algorithm.

There were several observations that stand out from the original experiments.

- All ways of using the algorithm in all cases reduced the amount of background noise.
- With too few coefficients, all that was heard were incoherent musical tones.
- As more coefficients were used, the sounds became more coherent but a growing amount of background noise crept in.
- Beyond a certain point, the addition of more coefficients contributed more and more to the noise volume while adding less and less to the speech volume.
- The gain from the extra levels was minimal. From the subjective point of view of the listener, there was little change after the first level and no improvement was noticeable after the second level.
- Prior knowledge of the sentence made a huge difference in the perception of the listener.
- The vowels were easy to identify and the fricatives disappeared.

14.1.3 The Wavelet Packet for Windows

The Wavelet Packet for windows, developed at Yale, allowed for a great deal more flexibility. It allowed for signals of up to 4096 samples in length; it offered up a wide collection of wavelets with some giving closer approximations of the original signal than others. It included software for the use of local cosines; it allowed for a number of ways to choose coefficients; the windows environment was a lot more user friendly.

With the increased flexibility came some downside. The 4,096 sample limit meant that to reconstruct an entire sentence with as many as 20,000 samples, several smaller intervals of speech had to be concatenated. This added a great deal of time to the process of experimenting.

There were some useful observations from the WPWL.

- The wavelets that came from scaling functions with longer recurrence relationships gave better results.
- For a given processing of the signal, the WPWL gave a good way of finding the approximately optimal choice of the number of coefficients, suggesting that algorithms for doing this could be built into a denoising algorithm.
- It was necessary for us to write our own programs to have the flexibility to deal with all the complications that would arise in developing an algorithm to meet the challenges posed by the problem of noisy sound signals.

14.2 Theory

In this section, our wavelet-based approach to noise reduction is described. This approach assumes that the noise is autoregressive (AR) and Gaussian, and uses the WPT of the noise's whitening filter to construct a second set of orthogonal basis vectors which are better suited for noise reduction. In this sense, the approach generalizes previous wavelet-based methods which assume the presence of white Gaussian noise (Saito [60]; Coifman and Donoho [12]; Lang et al. [44]; Pesquet et al. [53]).

14.2.1 The Wavelet Packet Transform

A typical implementation of the wavelet packet transform is shown in Fig. 14.1. There, a signal $\mathbf{x} = [x_0, x_1, \dots, x_{N-1}]$ (assumed periodic with period N) is input to a tree-structured bank of filter/decimator operations which divide the signal (in an approximate sense) into progressively narrower frequency sub-bands. The filters $H = \{h_{-n}\}_{n=0}^{M-1}$ shown in Fig. 14.1 are time reversed versions of the perfect-reconstruction conjugate quadrature filters derived by Daubechies [17]. The outputs of each sub-band's filter/decimator correspond to expansion coefficients $c_{j,k,m}$ for specific basis functions $\phi_{j,k,m}$, where $\{j, k, m\}$ approximately specify the time/frequency localization, frequency sub-band location, and temporal origin of the function. Basis functions for larger j tend, in an approximate sense, to be more precisely localized in the frequency domain and less precisely localized in the time domain. The transform coefficients for particular values of j are given as

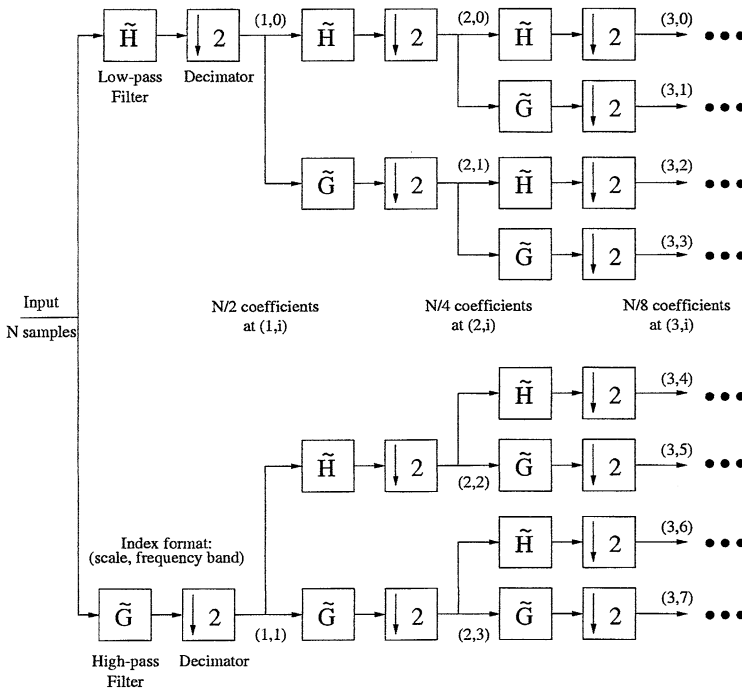


Fig. 14.1 Implementation of the wavelet packet transform

$$c_{j,k,m} = \begin{cases} \sum_{n=0}^{M-1} h_n c_{j-1,k/2,n+2m} & k \text{ even} \\ \sum_{n=0}^{M-1} g_n c_{j-1,(k-1)/2,n+2m} & k \text{ odd} \end{cases} \quad (14.1)$$

where $j \in \{0, 1, \dots, \log_2 N\}$, $k \in \{0, 1, \dots, 2^j - 1\}$, $m \in \{0, 1, \dots, 2^{-j}N - 1\}$, and $c_{0,0,m} = x_m$. Since the perfect-reconstruction property of the filters forces (14.1) to be a unitary transformation, the basis functions at any scale $j = J \geq 0$ can be constructed from linear combinations of basis functions at scales $j > J$. Hence, the collection of basis functions at all scales is redundant, and any collection of N basis functions corresponding to a disjoint covering of nodes of the tree provides a complete orthonormal basis for the signal.

The decimators used in Fig. 14.1 cause the WPT to be shift-invariant, such that applying a time-shift to \mathbf{x} can produce substantial changes in coefficient values. This shortcoming has led to the use of a shift-invariant WPT (or SIWPT), which implements the structure of Fig. 14.1 in a pre-selected basis without the decimators. The outputs of each filter then become the expansion coefficients in the pre-selected wavelet-packet basis for all circularly shifted versions of the input signal. The inverse SIWPT recombines these coefficients in a manner that effectively

reshifts the shifted signals to their origins and averages them together. More detailed discussions of shift-invariant transforms and their use in noise reduction are provided by Coifman and Donoho [12] and Pesquet et al. [53].

14.2.2 Selection of Wavelet Packet Bases

The speech estimates produced by subspace algorithms are, essentially, truncated series representations of the original noisy signal. It is well known that, of all orthonormal basis representations of signals in white noise, principal component bases produce the truncated series representations with the smallest mean-squared-error, and best compress the signal energy into the fewest number of coefficients. Watanabe [72] has shown that achieving these properties is equivalent to minimizing the “entropy” function

$$H(\{p_i\}) = - \sum_{i=0}^{N-1} p_i \log_2 p_i \quad (14.2)$$

where the “probability” p_i of the transform coefficient c_i is defined by

$$p_i = \frac{|c_i|^2}{\sum_{i=0}^{N-1} |c_i|^2}. \quad (14.3)$$

Wavelet packets, unlike principal-components, provide a multiplicity of bases from which to construct a “signal” subspace. Of these bases, the minimum-entropy basis provides the best signal compression. In previous studies assuming the presence of white noise, this basis was selected by means of the “best-basis” algorithm of Coifman and Wickerhauser [16], which used dynamic programming to select the wavelet-packet basis with minimum entropy and optimum signal compression. For the case of correlated noise, a generalized “best-basis” algorithm (Coifman and Saito [15]) is used which maximizes the “discrimination” function

$$H(\{p_i, q_i\}) = \sum_{i=0}^{N-1} p_i \log_2 (p_i/q_i) \quad (14.4)$$

where $\{p_i\}$ and $\{q_i\}$ respectively represent the “probabilities” of the noise-corrupted speech and the noise alone. This new criterion (which reduces to the entropy criterion in the case of white noise) provides a “local discrimination basis” (LDB) which gives the best possible separation between the “signal” and “noise” components of \mathbf{x} . The reader is referred to Coifman and Saito [15] for a more detailed explanation of this approach.

14.2.3 Separation of Signal and Noise Subspaces

After the optimum basis is selected for denoising, the basis functions must be separated into “signal” and “noise” subspaces. Previous studies have compared transform coefficients with a threshold to determine in which subspace their respective basis vector resides. Here, as in two previous studies (Saito [60]; Pesquet et al. [53]), the Minimum Description Length (MDL) model selection criterion of Rissanen [58] is used to drive the thresholds.

In the MDL framework for noise reduction, a noisy observation $\mathbf{x} \in \mathbb{R}^N$ (consisting of a signal \mathbf{s} in additive noise \mathbf{n}) is described by a theoretical binary codeword with length (in bits)

$$L(\mathbf{x}, \boldsymbol{\lambda}^{(k)}) = L(\boldsymbol{\lambda}^{(k)}) + L(\mathbf{x}|\boldsymbol{\lambda}^{(k)}) \quad (14.5)$$

$L(\boldsymbol{\lambda}^{(k)})$ and $L(\mathbf{x}|\boldsymbol{\lambda}^{(k)})$ are the lengths of codewords, respectively, describing $\boldsymbol{\lambda}^{(k)}$, a k^{th} order parametric model of \mathbf{x} , and the prediction error for the estimate $\hat{\mathbf{x}}(\boldsymbol{\lambda}^{(k)})$ derived from the parametric model. $L(\boldsymbol{\lambda}^{(k)})$ and $L(\mathbf{x}|\boldsymbol{\lambda}^{(k)})$ are, respectively, determined by a universal prefix coding method proposed by Rissanen and by the Shannon coding method. Among admissible parametric models, the model which produces the minimum description length is selected as the model most representative of the signal. In Saito [60] and Pesquet et al. [53], the parameter vector $\boldsymbol{\lambda}^{(k)}$ contained k transform coefficients and $N - k$ zeros, denoting respective “signal” and “noise” subspace contributions to \mathbf{s} . For this representation, Saito showed that

$$\boldsymbol{\lambda}^{(k)} = \left\{ \boldsymbol{\lambda} : \max_{\boldsymbol{\lambda} \in \Lambda(k)} \log_2 p_n(\mathbf{x} - \boldsymbol{\Phi}\boldsymbol{\lambda}) \right\} \quad (14.6)$$

$$L(\boldsymbol{\lambda}^{(k)}) = \frac{3k}{2} \log_2 N + C_1 \quad (14.7)$$

and

$$L(\mathbf{x}|\boldsymbol{\lambda}^{(k)}) = -\log_2 p_n(\hat{\mathbf{n}}(k)), \quad (14.8)$$

where $p_n(\mathbf{n})$ was the probability density of the noise, $\boldsymbol{\Phi}$ was an orthogonal matrix whose columns $\{\phi_i\}_{i=1}^N$ were the minimum-entropy wavelet packet basis for \mathbf{x} , $\Lambda(k)$ was the subset of vectors in \mathbb{R}^N with $N - k$ zero elements, $\hat{\mathbf{n}}(k) = \mathbf{x} - \boldsymbol{\Phi}\boldsymbol{\lambda}^{(k)}$ was the implicit noise estimate corresponding to the k^{th} -order signal model and C_1 was a constant independent of basis or model order.

In this study, the N elements $\{n_\ell\}_{\ell=0}^{N-1}$ of \mathbf{n} are assumed to be generated by an autoregressive (AR) process of the form

$$n_\ell = -\sum_{m=1}^p \alpha_m n_{\ell-m} + u_\ell \quad (14.9)$$

where $p \ll N$, and $\{u_\ell\}$ is a Gaussian white noise process with mean zero and variance of σ_u^2 . It has been shown (Gray [31]) that, when $p \ll N$, the linear convolution

in (14.8) may be approximated by a circular convolution, such that for (14.9) $\mathbf{u} = [u_0, u_1, \dots, u_n]^H$

$$\mathbf{A}^H \mathbf{n} = \mathbf{u} \quad (14.10)$$

where

$$\mathbf{A}^H = \begin{bmatrix} 1 & 0 & \dots & \alpha_p & \dots & \alpha_2 & \alpha_1 \\ \alpha_1 & 1 & 0 & \dots & \alpha_p & \dots & \alpha_2 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \alpha_p & \alpha_{p-1} & \dots & 1 & 0 & \dots & 0 \\ 0 & \alpha_p & \dots & \alpha_1 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \dots & \alpha_2 & \alpha_1 & 1 \end{bmatrix} \quad (14.11)$$

The probability density of \mathbf{n} is then

$$p_n(\mathbf{n}) = ((2\pi\sigma_u^2)^N \det(\mathbf{A}\mathbf{A}^H)^{-1})^{-\frac{1}{2}} \exp \left[-\frac{\mathbf{n}^H \mathbf{A}\mathbf{A}^H \mathbf{n}}{2\sigma_u^2} \right] \quad (14.12)$$

The codeword length for the prediction error of the k th-order estimate of the signal is then given by

$$\begin{aligned} L(\mathbf{x}|\boldsymbol{\lambda}^{(k)}) &= \frac{N \ln(2\pi\sigma_u^2) + \ln \det(\mathbf{A}\mathbf{A}^H)^{-1} + \mathbf{n}^H \mathbf{A}\mathbf{A}^H \mathbf{n} / \sigma_u^2}{2 \ln 2} \\ &= \frac{1}{2\sigma_u^2 \ln 2} \|\mathbf{A}^H \mathbf{n}\|^2 + C_2, \end{aligned} \quad (14.13)$$

where C_2 is a constant independent of basis or model order. After substituting $\hat{\mathbf{n}}(\boldsymbol{\lambda}^{(k)})$ for \mathbf{n} , $L(\mathbf{x}|\boldsymbol{\lambda}^{(k)})$ becomes

$$L(\mathbf{x}|\boldsymbol{\lambda}^{(k)}) = \frac{\|\mathbf{A}^H \boldsymbol{\Phi}_S(k) \boldsymbol{\lambda}_S^{(k)} + \mathbf{A}^H \boldsymbol{\Phi}_\eta(k) \boldsymbol{\lambda}_\eta^{(k)}\|^2}{2\sigma_u^2 \ln 2} + C_2 \quad (14.14)$$

where

$$\boldsymbol{\Phi}_S(k) = [\phi_{z_1}, \phi_{z_2}, \dots, \phi_{z_k}], \quad (14.15)$$

and

$$\boldsymbol{\Phi}_\eta(k) = [\phi_{z_{k+1}}, \phi_{z_{k+2}}, \dots, \phi_{z_N}]. \quad (14.16)$$

The index set $\{z_1, z_2, \dots, z_N\}$ is an ordering of the basis functions in which the first k basis functions are assumed to contain the signal and the last $N - k$ basis functions are assumed to contain only noise. The manner in which this partition is calculated is described below. The expressions

$$\boldsymbol{\lambda}_S^{(k)} = \boldsymbol{\Phi}_S(k)^H \hat{\mathbf{n}}(\boldsymbol{\lambda}^{(k)}), \quad (14.17)$$

and

$$\boldsymbol{\lambda}_\eta^{(k)} = \boldsymbol{\Phi}_\eta(k)^H \hat{\mathbf{n}}(\boldsymbol{\lambda}^{(k)}) \quad (14.18)$$

represent the projection of noise energy onto the “signal” and “noise” subspaces. Since the subspace spanned by $\{\phi_{z_i}\}_{i=k+1}^N$ is assumed to only contain noise, $\boldsymbol{\lambda}_\eta^{(k)} = \boldsymbol{\Phi}_\eta(k)^H \mathbf{x}$. The minimizing value of $\boldsymbol{\lambda}_s^{(k)}$ may then be found by solving a set of normal equations, leading minimization of $L(\mathbf{x}|\boldsymbol{\lambda}^{(k)})$ to become equivalent to maximization of $\|\mathbf{A}^H \boldsymbol{\Phi}_s(k) \boldsymbol{\Phi}_s^H(k) \mathbf{x}\|^2$. \mathbf{A}^H , which contains N circular shifts of the prediction error filter, can be written as

$$\mathbf{A}^H = \mathbf{F} \mathbf{Q}_A \mathbf{F}^H \quad (14.19)$$

where $[\mathbf{F}]_{mn} = \frac{e^{j2\pi mn/N}}{\sqrt{N}}$, and \mathbf{Q}_A is a diagonal matrix containing the value of the FFT of the first row of \mathbf{A}^H divided by \sqrt{N} . Sinha and Tewfik ([64], p.3476) have shown that, for the Daubechies filters having $M \geq 10$, $\boldsymbol{\Phi}^H \mathbf{F} \mathbf{Q}_A^H \mathbf{Q}_A \mathbf{F}^H \boldsymbol{\Phi}$ is nearly diagonal. Hence, the columns of $\mathbf{A}^H \boldsymbol{\Phi}_s(k)$ are nearly orthogonal,

$$\|\mathbf{A}^H \boldsymbol{\Phi}_s(k) \boldsymbol{\Phi}_s^H(k) \mathbf{x}\|^2 \approx \sum_{i=1}^k (\phi_{z_i}^H \mathbf{x})^2 \|\mathbf{A}^H \phi_{z_i}\|^2, \quad (14.20)$$

and

$$L(\mathbf{x}|\boldsymbol{\lambda}^{(k)}) \approx \frac{\|\mathbf{A}^H \mathbf{x}\|^2 - \sum_{i=1}^k (\phi_{z_i}^H \mathbf{x})^2 \|\mathbf{A}^H \phi_{z_i}\|^2}{2\sigma_u^2} + C_2 \quad (14.21)$$

where $\{z_i\}_{i=1}^k$ are the indices of the set of filtered basis vectors with the largest coefficients, and the values $\{\mathbf{A}^H \phi_{z_i}\}$ are computed from the SIWPT of a time reversed version of the whitening filter $\{\alpha_i\}$. By taking the finite difference (with respect to k) of the sum of (14.7) and (14.21), and looking for the value of k for which that difference changes from negative to positive, the MDL criterion may be shown to select the largest value of k for which

$$|\phi_{z_k}^H \mathbf{x}| > \frac{\sigma_u}{\|\mathbf{A}^H \phi_{z_k}\|} \sqrt{3 \ln N}. \quad (14.22)$$

Hence, for subspace denoising using a single shift-invariant WPT, evaluation of $L(\mathbf{x}|\boldsymbol{\lambda}^{(k)})$ for various values of k may then be replaced by simple threshold comparisons for each basis vector.

In this study, as in every other study of subspace methods, truncating the expansion of the noisy signal imposes some audible artifacts on the enhanced speech. These artifacts may be reduced by averaging together several denoised versions of the signal as described in Sect. 14.2.1 with the SIWPT. Note that each SIWPT coefficient for a given scale j and sub-band k corresponds to a shift-variant WPT

coefficient in more than one circularly shifted version of the signal. Therefore, averaging N denoised, circularly shifted versions of the signal requires that each coefficient be subjected to multiple threshold rules.

14.2.4 Summary of the Algorithm

This section has described a wavelet based approach for enhancing signals in additive autoregressive noise. The algorithm uses the SIWPT of the noise's prediction-error filter to implement a form of transform-domain convolution with wavelet packet basis vectors. The sequence of operations in this algorithm is summarized as follows:

1. Use conventional methods to obtain an AR model of the noise.
2. Derive coefficient variance estimates for both the noise and the noise corrupted speech.
3. Use the variance estimates to select the LDB as described above and in Coifman and Saito [15].
4. Compute the SIWPT of the signal and the filter response in the LDB,
5. Use (14.21) to select coefficients assumed to represent the signal, and attenuate noise coefficients assumed to represent noise.
6. Resynthesize the signal by means of the inverse SIWPT.

In the next section, we will present experimental results which demonstrate the utility of the proposed algorithm.

14.3 Objective Evaluation

This experiment compares the described approach with a conventional spectral subtraction algorithm (Schwarz et al. [63]).

14.3.1 Stimuli

A recording of the Harvard/IEEE [35] sentence “That hose can wash her feet” (followed by 800 msec of silence) was used for all trials of the experiment. (A spectrogram of the signal is shown in Fig. 14.2).

The recording was sampled at 8 kHz and combined separately with two different noise maskers to produce two sets of waveforms with overall SNRs of 0, 5, 10, 15, and 20 dB. The first noise masker was derived from a 12 s segment of digitally recorded noise in an automobile traveling at highway speed. The second noise




Fig. 14.2 “That hose can wash her feet”




Fig. 14.3 Spectra of simulated car noise

masker was a sequence of autoregressive gaussian noise, produced by a fourth order all pole filter of the form

$$n_k = 1.3733n_{k-1} - 0.7858n_{k-2} + 0.1138n_{k-3} + 0.0119n_{k-4} + u_k \quad (14.23)$$

where $\{u_k\}$ was a sequence of white Gaussian noise, and the filter parameters were derived from a 2 s segment of digitally recorded car noise. The AR model produces a noise spectrum which is relatively flat below 1 kHz, and rolls off at about 12 dB/octave above 1 kHz, as shown in Fig. 14.3. A spectrogram of the noise corrupted sentence at 5 dB SNR is shown in Fig. 14.4.

Noise maskers of this type were selected for two reasons. First, the car noise provides an excellent example of a real-world masker that can be modeled suitably by autoregressive noise. Second, the car noise’s inherent non-stationarity allows us to evaluate the algorithm’s performance in settings which approximate practical applications.




Fig. 14.4 Sentence of Fig. 14.2, in simulated car noise at 5 dB SNR

The ten noise-corrupted speech signals were processed in 50% overlapped segments by a version of the SIWPT algorithm that uses a single-frame estimate of the noise's whitening filter. In previous work (Whitmal et al. [74]), that algorithm improved the SNR of speech signals (taken in 256 sample frames) in simulated car noise at 0 dB by 11.68 dB, and reduced noise in pause regions by as much as 23.77 dB. The signals were also processed by a spectral subtraction algorithm similar to that of Schwartz et al. [63], which estimates the signal by subtracting a weighted version of an estimated noise power spectrum from the observed signal's power spectrum. Frame lengths of 64 msecs (512 samples at 8 kHz) were used in both algorithms to produce enhanced signals with low residual noise output. For reference, signals were also processed by a SIWPT algorithm (similar to that of Pesquet et al. [53]) designed to reduce white noise.

Both Schwartz et al. [63] and Boll [4] note that the spectral subtraction algorithm's parameters must be carefully selected if the algorithm is to provide a good tradeoff between signal distortion and residual noise output. Preliminary listening indicated that a noise floor constant of .01 and oversubtraction by a factor of 5 would produce intelligible speech with acceptable amounts of residual noise. The long frame length contravenes the guidelines of Schwartz et al. [63]. In informal listening, however, spectral subtraction output for a 64 msec frame length sounded less noisy and more faithful than output using a 16 or 32 msec frame length. The reader is referred to Schwartz et al. [63] for a more detailed discussion of these parameters and their effects on noise reduction.

14.3.2 Methods

The performance of the algorithms was evaluated by means of two objective measures: segmental SNR and log spectral distance. These measures are respectively defined as

$$SNR = \frac{1}{L} \sum_{i=1}^L 20 \log_{10} \frac{\|\mathbf{s}_i\|^2}{\|\mathbf{s}_i - \hat{\mathbf{s}}_i\|^2} \quad (14.24)$$

and

$$LSD = \frac{1}{NL} \sum_{i=1}^L \left[\sum_{m=0}^{N-1} \left| 20 \log_{10} \frac{|S_i(m)|}{|\hat{S}_i(m)|} \right| \right], \tag{14.25}$$

where s_i and \hat{s}_i are the actual and estimated speech signals in the i th of L N -sample frames of s and $S_i(\cdot)$ and $\hat{S}_i(\cdot)$ are their respective N -sample DFTs. These measures were selected both because of their minimal computation requirements, and because of their high correlations with subjective speech quality measures (e.g., the Diagnostic Acceptability Measure (DAM) of Voiers [71]). The reader is referred to Quackenbush et al. [56] for a detailed discussion of the use of these measures.

Measurements of segmental SNR and log spectral distance were computed for non-overlapping 256-sample segments and averaged over speech regions to produce objective measures of quality. Data for both measures are shown in Tables 14.1 and 14.2. To assess noise reduction in the silence regions, the RMS level of residual noise was computed and compared with the RMS level of the original noise masker. Comparisons of these noise levels are presented in Table 14.3.

Table 14.1 Segmental SNR measures (in dB) for spectral subtraction and SIWPT-based subspace denoising algorithms in two types of noise

SNR	Control		SIWPT, prop.		SIWPT, conv.		Spec. sub.	
	S	R	S	R	S	R	S	R
0	−10.21	−10.30	1.92	−0.20	−3.34	−4.76	−2.01	−1.88
5	−5.20	−5.30	4.53	2.98	0.39	−0.85	1.82	1.94
10	−0.20	−0.30	7.27	6.26	4.00	2.95	5.32	5.55
15	4.80	4.70	10.12	9.37	7.60	6.70	8.47	8.86
20	9.80	9.71	13.55	12.77	11.27	10.70	11.38	11.80

In each of the tables, measures taken in simulated and real car noise are respectively denoted by “S” and “R” where R = real automobile road noise and S = synthesized road noise. Parenthesized figures indicate number of frames (out of 34) correctly estimated to contain silence

Table 14.2 Noise reduction measures (in dB) for spectral subtraction and SIWPT-based subspace denoising algorithms in two types of noise

SNR	SIWPT, prop.		SIWPT, conv.		Spec. sub.	
	S	R	S	R	S	R
0	42.55 (19)	18.49 (2)	11.34 (0)	8.05 (0)	12.96 (0)	12.77 (0)
5	42.72 (21)	18.49 (4)	11.35 (0)	8.06 (0)	12.96 (0)	12.77 (0)
10	42.89 (25)	18.50 (4)	11.37 (0)	8.06 (0)	12.96 (0)	12.77 (0)
15	43.15 (26)	18.51 (6)	11.39 (0)	8.07 (0)	12.96 (0)	12.77 (0)
20	43.54 (29)	18.56 (6)	11.38 (0)	8.10 (0)	12.96 (0)	12.77 (0)

R = real automobile road noise
S = synthesized road noise

Table 14.3 Log spectral distance measures (in dB) for spectral subtraction and SIWPT-based subspace denoising algorithms in two types of noise

SNR	Control		SIWPT, prop.		SIWPT, conv.		Spec. sub.	
	S	R	S	R	S	R	S	R
0	18.16	17.87	39.27	15.82	14.01	13.77	11.83	11.60
5	14.71	14.58	32.91	12.62	12.97	12.87	10.13	9.77
10	11.70	11.54	24.96	11.19	11.37	12.61	8.50	8.09
15	9.20	9.03	21.61	9.37	10.74	11.07	8.47	8.86
20	7.00	6.85	17.50	7.92	9.99	9.86	5.59	5.47

R = real automobile road noise
S = synthesized road noise




Fig. 14.5 Sentence of Fig. 14.4 processed by proposed subspace algorithm




Fig. 14.6 Sentence of Fig. 14.4, processed by spectral subtraction

14.3.3 Results

The data of Table 14.1 and Fig. 14.5 indicate that, for both types of noise, the SIWPT algorithm described above improves segmental SNR better than the standard spectral subtraction algorithm and the conventional SIWPT algorithm. These results are particularly encouraging for the case of actual car noise, where, as expected, the proposed algorithm’s performance decreases. Spectrograms of sentences processed by the SIWPT algorithm and by the spectral subtraction algorithm in simulated noise are shown in Figs. 14.5 and 14.6. The presence of spectral subtraction’s “musical noise” is indicated by the large number of short-duration, high-energy bursts visible in Fig. 14.6. For reference, a spectrogram of output from the conventional algorithm is shown in Fig. 14.7




Fig. 14.7 Sentence of Fig. 14.4, processed by conventional subspace algorithm




Fig. 14.8 Log spectral distance for the SIWPT method at 0 dB SNR

In contrast, the data of Table 14.3 indicate that the spectral subtraction algorithm improves log spectral distance better than the proposed algorithm, particularly in simulated car noise. The increase in distance is largely due to the algorithm’s strong attenuation of low energy consonants which often resemble gaussian noise. This phenomenon is illustrated in Fig. 14.8, which compares log spectral distance measures for the proposed algorithm in both types of noise at 0 dB SNR. As indicated by the figure, distance measures for the two noise types are close in value in regions of moderate SNR. The conflicting trends in the two objective measures provide further illustration of the tradeoffs made between noise reduction and signal distortion. In understanding this conflict, it should be noted that the correlations between segmental SNR and DAM scores demonstrated in Quackenbush et al. [56] have

only been confirmed for waveform coding methods. It is not clear how the type of distortion demonstrated in Fig. 14.8 will affect segmental SNR measurements, since the present approach bears a closer resemblance to transform coding.

14.4 Subject Testing

The proposed noise reduction algorithm was tested on hearing impaired individuals at the Northwestern University Department of Communication Disorders.

14.4.1 Subjects

Participants in this study consisted of 16 Northwestern University Hearing Clinic patients with impaired hearing acuity. Eligible patients were screened on the basis of their audiograms and assigned to one of two categories: “flat” (i.e., hearing loss slope above 500 Hz \leq 10 dB/octave) and “sloped high-frequency” (i.e., hearing loss above 500 Hz $>$ 10 dB/octave). Each category consisted of eight subjects. Mean hearing losses for each group are shown in Table 14.4.

14.4.2 Test Protocol

The new method was evaluated in trials of the Revised Speech Perception in Noise (R-SPIN) test [3] at the facilities of the Northwestern University Hearing Clinic in Evanston, Illinois. The R-SPIN test protocol uses eight lists of 50 recorded sentences, each of which are presented to subjects in the presence of interfering noise. Subjects attempt to identify the last word of each sentence. In each list, the last word of 25 of the sentences is predictable from context (e.g., “Stir your coffee with a spoon”), while the last word of the other 25 sentences is difficult to predict from context (e.g., “Bob could have known about the spoon”). These two types of sentences are denoted as “PH” and “PL” sentences. This approach allows researchers to simultaneously assess how well subjects are able to both hear and interpret noise-corrupted speech.

Recordings of R-SPIN sentences were sampled at 8 kHz and combined electronically with digital recordings of automobile road noise at signal-to-noise ratios (SNRs) of 5 dB and 10 dB. Individual tapes of the eight lists were then made for each subject participating in the test. For every tape, each of the eight lists was assigned one of eight possible combinations of the two SNR levels and one of four processing options:

1. Control (sentence in noise, no processing).
2. Linear filtering per the Australian National Acoustic Laboratories’ (NAL) hearing-aid gain-prescription rule [9].

Table 14.4 Profiles of impaired-acuity subjects

Flat-loss subjects								
Demographics				Hearing loss (dB HL)				
Subject	Age	Sex	Ear	250 Hz	500 Hz	1 kHz	2 kHz	4 kHz
FC	77	F	R	20	30	45	55	55
BF	74	M	R	45	45	50	55	60
RF	61	M	R	60	55	70	65	80
IG	86	M	L	40	50	45	45	65
DK	18	F	R	45	45	50	55	50
LL	82	F	R	20	20	25	40	35
CM	77	M	L	25	30	45	50	60
MM	86	M	L	15	25	40	40	40
Means	70			34	40	46	49	56
SDs				15	12	12	8	13

Sloped-loss subjects								
Demographics				Hearing loss (dB HL)				
Subject	Age	Sex	Ear	250 Hz	500 Hz	1 kHz	2 kHz	4 kHz
FA	56	M	L	35	30	40	50	55
LK	76	M	R	25	30	25	40	60
RL	66	M	L	25	25	30	55	65
DM	82	F	L	30	25	40	50	75
SM	88	F	R	15	20	20	50	65
HM	83	M	L	10	30	50	55	55
HN	83	M	R	20	25	30	50	65
HP	71	M	R	20	25	45	55	50
Means	76			23	28	35	51	60
SDs				8	3	10	5	7

Standard deviations are rounded to nearest dB to maintain format

- 3. The shift-invariant wavelet-packet transform (SIWPT) based noise reduction algorithm of [74].
- 4. A combination of NAL filtering and SIWPT-based processing.

The NAL gain-prescription rule is a formula (derived from clinical data) which assigns approximately .36 dB of gain for every 1.0 dB of hearing loss. NAL prescriptions for each patient were pre-calculated from previous audiological measurements, and used to design conventional FIR digital filters for convolution with R-SPIN sentence recordings. RMS levels for each sentence were then matched to that of a calibration tone before being transferred to digital audio tape. Each tape, which contained a training list of R-SPIN sentences in quiet followed by the processed sentences, was then presented monaurally over headphones to the subject at that subject’s most comfortable listening level.

14.4.3 Experimental Results

Mean R-SPIN scores for both the flat and sloped subject groups are presented in Table 14.5. Inspection of these scores reveals several interesting trends. In particular, for flat-loss subjects, NAL filtering appeared to have a negative effect on many R-SPIN scores. The proposed SIWPT method, in contrast, appeared to improve R-SPIN scores for flat-loss subjects.

A repeated-measures analysis of variance (ANOVA) was conducted to evaluate the dependence of R-SPIN scores on impairment type, subject, sentence context, SNR, and processing type. The results of the ANOVA indicate that the factors of subject, context and SNR have statistically significant ($p < .001$) effects on R-SPIN scores. Context effects on R-SPIN scores are consistent with effects seen in previous research [8]. However, the results of the ANOVA do not indicate the differences in processing have a significant effect on R-SPIN scores.

The significance of subjects as a factor led us to conduct an analysis of individual R-SPIN scores. The Thornton–Raffin binomial model of test scores [68] was used to construct 95% confidence intervals for each of the control scores. Scores for the non-control process were then compared to the confidence interval limits of the corresponding control and subsequently classified as “better” , “same” or “worse” than the control score. The results of these comparisons (presented in Table 14.6) suggest that SIWPT processing tends to improve R-SPIN scores at 10 dB SNR, and

Table 14.5 Mean R-SPIN scores for subjects with impaired acuity

Context/SNR	Processing for flat-loss subjects			
	Control	NAL	SIWPT	NAL+SIWPT
PH, 5 dB	18.63	16.00 ^a	18.00	14.75 ^a
PH, 10 dB	22.13	20.75	23.13	22.38
PL, 5 dB	7.25	7.38	8.13	8.88
PL, 10 dB	10.50	10.75	13.13 ^a	12.50
Context/SNR	Processing for sloped-loss subjects			
	Control	NAL	SIWPT	NAL+SIWPT
PH, 5 dB	18.13	16.13	16.13	15.88
PH, 10 dB	21.75	22.25	21.25	21.75
PL, 5 dB	7.75	6.63	7.63	6.63
PL, 10 dB	10.13	10.38	10.00	9.75

^adenotes significance per Dunnett T-test at ~0.06 level

Table 14.6 Distribution of relative R-SPIN score ratings

Rating	5 dB SNR			10 dB SNR		
	Worse	Same	Better	Worse	Same	Better
NAL	3	11	2	2	12	2
SIWPT	3	11	2	1	11	4
Both	3	12	1	2	12	2

that both types of processing tend to reduce R-SPIN scores at 5 dB SNR. These results appear to be strongly influenced by the flat loss group, which accounted for six of the nine subjects reporting at least one significant deviation.

14.5 Built-in Distortions

Both the original and proposed subspace algorithms construct “signal” subspaces from the subset of basis vectors which produce the largest transform coefficients. The conventional algorithm uses a standard wavelet-packet basis (where all basis vectors contain the same amount of energy), and selects basis vectors solely on the basis of coefficient size. Since both the original speech and the car noise have most of their energy at low frequencies, the conventional algorithm tends to preserve basis vectors with substantial energy at low frequencies. In contrast, the filtering imposed on the new algorithm’s basis vectors increases the energy in vectors whose spectrum is most severely compromised by the noise, thereby offsetting the low SNRs in these subspaces.

Some comments about the objective data are appropriate at this point. First, we note that the combination of coefficient attenuation and averaging can produce benign distortion artifacts which result in objective measures with overly pessimistic values. (This is particularly true at the 20 dB SNR). Of greater concern is the tendency of this algorithm to attenuate both speech and noise in regions of low SNR. This tradeoff is particularly noticeable in certain consonant regions, as indicated by Figs. 14.2, 14.4, and 14.5. This tendency is a direct consequence of the binary threshold measure used in wavelet-packet based denoising algorithms. Use of this approach is motivated by two factors. First, use of the entire coefficient value (rather than a scaled version) appears to provide smaller values of mean-squared error (Coifman and Donoho [12]). Second, in algorithms using wavelet-packet transforms with “best basis” approaches, the popular linearly-tapered thresholds used in many wavelet transform based studies (e.g., see Donoho [18]) appear to have little effect on the signal. This may be explained heuristically by noting that the “best-basis” algorithm increases the size of “signal” and reduces the size of the “noise” coefficients to the point where thresholding effects on either type of coefficient are reduced considerably. It is possible that improved performance may be obtained by using several Bayesian threshold approaches (Vidakovic [69]; Chipman et al. [11]) which were developed in concurrence with the present work.

The results of Sect. 14.4.3 demonstrated that the proposed method fails to provide significant improvements in intelligibility despite improvements in objective measures. One possible explanation is the level-dependent nature of the attenuation that the algorithms impose on the speech. In the following section, we show (both theoretically and empirically) that the proposed algorithm attenuates input speech in a manner that resembles the attenuation produced by recruitment-of-loudness simulators, such that their attenuation can be equated with an effective sensorineural hearing loss.

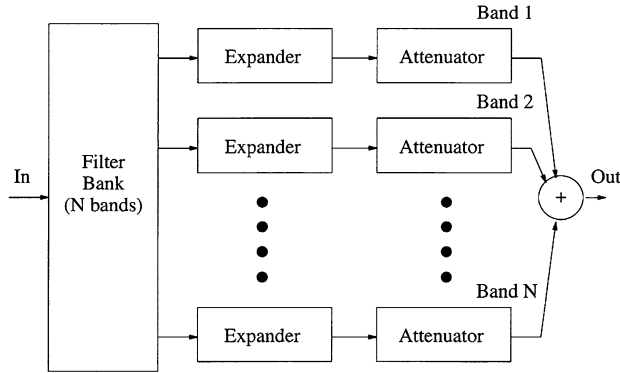


Fig. 14.9 Block diagram of ideal recruitment-of-loudness simulator

Figure 14.9 shows an idealized processor that simulates the effects of recruitment of loudness for normal-hearing listeners. The processor uses a filter bank to partition the signal into several frequency bands. Signals in each band are then input to a combination of amplitude expansion and attenuation modules. The expansion ratio of each expander is set equal to the ratio of the dynamic range in that band for a normal-hearing listener to the dynamic range in that band for the hearing-impaired listener. The attenuation factor is likewise determined such that the speech has the same relative intensity to the normal threshold as the corresponding element of the unprocessed speech has to the impaired threshold.

An example of recruitment simulation is provided by Duchnowski and Zurek [24], who used piecewise-linear expansion characteristics in each of 14 frequency bands. The input/output relationship in each band was given by the relation

$$y = \begin{cases} Kx + (1 - K)T_c, & x < T_c \\ x, & x \geq T_c \end{cases}$$

where x is the input level in dB SPL, (sound pressure level), y is the output level in dB SPL, T_c is the effective level of complete recruitment (i.e., the level at which loudness perception in the impaired ear equals that of the normal ear), and K , in turn, is given as

$$K = \frac{T_c - T_N}{T_c - T_I}$$

where T_N and T_I are (respectively) the normal and impaired thresholds of hearing, and $T_I - T_N$ is the hearing loss (HL) modeled by the simulator. An illustration of the expander characteristic is shown below in Fig. 14.10. Values of T_c were derived from the data of Hallpike and Hood [34], who expressed K as the tangent of a “recruitment angle” which was proportional to the degree of hearing loss. The data of Duchnowski and Zurek [24] indicate good agreement between intelligibility scores of subjects with moderate sensorineural hearing losses and normal-hearing subjects listening to the output of simulators matched to the subjects’ audiograms.




Fig. 14.10 Input-output characteristic of Duchnowski-Zurek recruitment simulator (14.26)

The attenuation of characteristics produced by these simulators are also produced by subspace algorithms like the proposed method. Subspace algorithms have typically employed two types of attenuation rules: “hard” threshold rules and “soft” threshold rules. Hard threshold rules apply no attenuation to “speech subspace” coefficients (i.e. all coefficients above the noise threshold), and infinite attenuation to “noise subspace” coefficients. An approximate expression for hard-thresholding effects can be derived for algorithms using orthonormal transforms, using common (Jansen [37]) assumptions of deterministic coefficients in additive white Gaussian noise. (Results derived for additive white noise can be extended to cases of colored noise through the use of whitening filters, as discussed in Ephraim and van Trees [28]). For a particular basis function, let the observed coefficient X (a sum of signal coefficient S and random noise coefficient D) be subjected to hard-thresholding. D is assumed to have a mean of zero and a variance of σ_D^2 . The threshold coefficient is given as

$$V = \begin{cases} X, & |X| > T \\ 0, & |X| \leq T \end{cases}.$$

V has the probability distribution

$$f_V(v; S) = \left[\left(\frac{1}{\sqrt{2\pi\sigma_D^2}} e^{-(v-S)^2/2\sigma_D^2} \right) (u(v-T) + u(-v-T)) \right] + \Pr(S)\delta(v),$$

where $u(v)$ is the Heaviside step function, $\delta(v)$ is the Dirac δ function and $\Pr(S)$, the probability that the coefficient S is attenuated, is

$$\Pr(S) = \frac{1}{2} \left[\operatorname{erf} \left(\frac{T+S}{\sigma_D \sqrt{2}} \right) + \operatorname{erf} \left(\frac{T-S}{\sigma_D \sqrt{2}} \right) \right].$$

Using integration by parts (Stein [66]), we can show that the expected energy in the transform coefficient is

$$\begin{aligned} E \{v^2\} &= \int_{-\infty}^{\infty} v^2 f_V(v) dv \\ &= (S^2 + \sigma_D^2) (1 - \Pr(S)) \\ &\quad + \frac{\sigma_D^2}{\sqrt{2\pi}} \left(\frac{T+S}{\sigma_D} \right) e^{-(T-S)^2/2\sigma_D^2} + \frac{\sigma_D^2}{\sqrt{2\pi}} \left(\frac{T-S}{\sigma_D} \right) e^{-(T+S)^2/2\sigma_D^2} \end{aligned}$$

or

$$\begin{aligned} E \{v^2\} &= \sigma_D^2 \left(\left(\frac{S}{\sigma_D} \right)^2 + 1 \right) (1 - \Pr(S)) \\ &\quad + \frac{\sigma_D^2}{\sqrt{2\pi}} \left[\left(\tau + \frac{S}{\sigma_D} \right) \right] e^{-(\tau-S)^2/2\sigma_D^2} + \frac{\sigma_D^2}{\sqrt{2\pi}} \left[\left(\tau - \frac{S}{\sigma_D} \right) \right] e^{-(\tau+S)^2/2\sigma_D^2} \end{aligned} \quad (14.26)$$

where $\tau = \frac{T}{\sigma_D}$. (Typically $\tau > 2\sigma_D$.) Setting $S_{\text{dB}} = 20 \log_{10} |S|$ and $V_{\text{dB}} = 10 \log_{10} (E \{v^2\})$ gives

$$\begin{aligned} V_{\text{dB}} &= 10 \log_{10} \sigma_D^2 + 10 \log_{10} \left[\left(\left(\frac{10^{S_{\text{dB}}}}{\sigma_D} \right)^2 + 1 \right) (1 - \Pr(10^{S_{\text{dB}}/20})) \right] \\ &\quad + \left(\tau + \frac{10^{S_{\text{dB}}/20}}{\sigma_D} \right) \frac{e^{-\left(\tau - \frac{10^{S_{\text{dB}}}}{20}\right)^2/2\sigma_D^2}}{\sqrt{2\pi}} + \left(\tau - \frac{10^{S_{\text{dB}}/20}}{\sigma_D} \right) \frac{e^{-\left(\tau + \frac{10^{S_{\text{dB}}}}{20}\right)^2/2\sigma_D^2}}{\sqrt{2\pi}} \end{aligned} \quad (14.27)$$

At high input levels, where $S \gg T$, $\Pr(S) \approx 0$ and $e^{-(\tau \pm S)^2/2\sigma_D^2} \approx 0$, it follows that $E \{v^2\} \approx 0$ and $V_{\text{dB}} = S_{\text{dB}}$. At input levels below T , (14.26) exhibits level-dependent attenuation (or equivalently, an accelerated rate of increase in output level) that resembles recruitment of loudness simulation. An example of this behavior (with all regions marked) is shown in Fig. 14.11, with $\sigma_D = 10$ and $\tau = 5$. The effects of the choice of τ are illustrated in Fig. 14.12, which plots input and output levels (normalized by σ_D) for $3 \leq \tau \leq 7$. Figure 14.12 shows that large values of τ produce steep recruitment slopes, and small values of τ produce shallow slopes. Hence, τ




Fig. 14.11 A sample recruitment curve for subspace noise reduction ($\sigma_D = 10, \tau = 5$)




Fig. 14.12 Effects of τ on subspace noise reduction performance

determines both the rate of attenuation and the level of residual noise. Note that the energy in each coefficient is typically spread over a band of frequencies, rather than being concentrated at a single frequency.

The idealized input–output relation of (14.27) is a smooth, differentiable function; as such, standard numerical algorithms can be used to find effective hearing

losses for any given output level. However, for the present algorithm, the signal and noise subspaces are time-invariant, and selected on the basis of instantaneous values of signal and noise. The time-varying nature of these algorithms makes it difficult to derive an expression for their expected coefficient values. One way to address this problem is to fit long-term measurements of input and output level to a parametric model that characterizes the algorithm's level-dependent attenuation, and can be used to determine the effective hearing loss imposed by the algorithm.

This task was performed by Whitmal and Vosoughi [77], who measured speech and noise levels for recordings of the syllables “of”, “ook”, “od”, “pa”, “ja”, and “ma”, taken (respectively) from word lists from the Nonsense Syllable Test (Levitt and Resnick, [48]). The syllables were specifically selected to give an appropriately diverse representation of the syllables used in the test. Models for the processed signals were derived from RMS level measurements taken in critical bands. Clean speech at 74 dB SPL was filtered into twelve bands centered at 100, 250, 500, 750, 1,000, 1,500, 2,000, 3,000, and 4,000 Hz. Since the non-linear expansion of the wavelet-based algorithm imposes smearing on the signal spectrum, prefiltering was used to isolate processed components in one band from leakage of processed components in other bands. Simulated automobile noise at 56 dB SPL was added electronically to each of the nine signals which were then processed by the wavelet-based noise reduction algorithm. (The use of broadband noise insures that the algorithm uses appropriate thresholds for each of the narrow-band signals.)

Following processing, the output signals were upsampled to 24,000 Hz and input to a bank of critical filters with center frequencies 100, 250, 500, 750, 1,000, 1,500, 2,000, 3,000, 4,000, 6,000, 8,000 and 10,000 Hz. The output of each filter was then passed through an RMS level detector (time constant: 21.3 msec) and converted to decibels. The RMS level measurements for the processed speech were combined in time-aligned ordered pairs with concurrent RMS level measurements for clean speech, plotted on scattergrams (see Fig. 14.14) and fitted to piecewise-linear functions from which values of T_N and T_I could be extracted. The resulting effective hearing loss is plotted in Fig. 14.13 according to audiological conventions, where normal hearing (0 dB loss) is at the top of the y-axis and profound loss (120 dB) at the bottom. The present algorithm produces an effective hearing loss for speech of about 30 dB between 500 and 3,000 Hz. In combination with a patient's hearing loss, the present algorithm would be expected to result in poorer intelligibility.

14.6 Conclusions

The presence of ambient noise presents a major problem for hearing impaired listeners. When amplified by a hearing aid, it creates an annoying sound that degrades the intelligibility of speech. The results of our experiments with noise reduction were mixed. As described in Sect. 14.3, our algorithm was quite successful in reducing noise; however, as reported in Sect. 14.4, it produced no comparable improvement in intelligibility when tested on hearing impaired subjects. The reduction of noise,




Fig. 14.13 Effective hearing loss due to hard-thresholding




Fig. 14.14 A recruitment curve for wavelet-packet subspace noise reduction measured at 500 Hz

in and of itself, is helpful because it improves the quality of sound. But people use hearing aids because they have trouble understanding conversations, particularly in noisy environments. Even people with normal hearing can have difficulty understanding a conversation in a crowded room with several conversations occurring simultaneously. Ideally, a hearing aid would be able to both reduce noise and improve intelligibility, a goal that has eluded hearing aid producers. In Sect. 14.5,

we give a partial explanation of what makes this problem so challenging. Whitmal and Vosoughi found [77] that the output of our algorithm is similar to a process that simulates recruitment of loudness (a hearing problem) for listeners with normal hearing.

Recruitment of loudness is a hearing disorder that is characterized by compression of a person's dynamic range. The term dynamic range refers to the level of loudness between the barely audible and the painfully loud. It can vary a great deal from one individual to another, and more importantly, from one frequency to another. With recruitment of loudness, the threshold of audibility is raised, usually more heavily in the higher frequencies, without a comparable rise in the threshold of pain. The resulting compression of the dynamic range causes distortions in the processing of speech signals leading to a loss of intelligibility by the impaired listener. Older people, for example, often find it difficult to understand the high pitched voices of children. What Whitmal and Vosoughi [77] showed was that hard-thresholding, a feature of our algorithm, produced effects that resembled recruitment of loudness.

Our experiments suggest a more fundamental problem. Algorithms that reconstruct a signal from a small percentage of the transform coefficients lose information. Data compression algorithms, for example, frequently result in some signal distortion. The television reception coming from uncompressed digital signals via an antenna compares favorably with the reception obtained from the compressed digital signals transmitted via cable. The full signal clearly produces a better image. Nevertheless, the amount of distortion is not very noticeable and most users find that the advantages of cable outweigh its disadvantages.

Improving the intelligibility of noisy speech is far more challenging. Speech signals contain some redundancy in the sense that, for example, if the high frequency components of recorded speech are filtered out, it may still be recognizable to a person with normal hearing. However, some of these redundancies are essential for hearing impaired listeners who may no longer be able to understand the filtered speech. Comprehending speech involves more than just the ear's processing of the signal. It involves the way the mind organizes and processes information. The loss of some of the speech signal that occurs with the denoising algorithm reduces the ability of the impaired listener to do so.

When we started our experiments with denoising algorithms, people who worked on hearing aids told us that the key to whether wavelet packets would help was in how well they could match the shape of speech signals. The linear discriminant basis procedure produced a "signal" subspace that is only a "most likely to be signal" subspace. In the end the algorithm did not project onto a basis of phonemes but projected onto a basis of wavelet packets that may have collectively resembled human speech, but was not close enough to render the noisy sentences more intelligible. The wavelet packets were built from a small library of wavelet functions. Though the number of possible wavelet packet bases generated by a single wavelet was quite numerous, the geometry of the different wavelet packet basis functions was similar. The underlying library of wavelets was simply not robust enough to match the components of speech.

Recent work in the area of the geometry of large data sets has given rise to a whole new generation of orthonormal bases. These bases are adapted to the data they are describing and so offer the promise of overcoming some of the intelligibility challenges in the construction of hearing aids. A special issue of the online journal *Applied and Computational Harmonic Analysis* devoted to diffusion maps [1] introduced novel ways of applying the tools of harmonic analysis to the study of the geometry of large data sets. The paper of Coifman and Maggioni, *Diffusion wavelets* [13], and the companion paper of Bremer, Coifman, Maggioni, and Szlam, *Diffusion wavelet packets* [7], provide a framework in which diffusion wavelet packets could be used to denoise speech signals in a manner similar to the research described in this paper. The geometry of diffusion wavelet packets, matching more closely the geometry of human speech, might be able to reduce the noise without losing as much of the signal.

Partha Niyogi and Aren Jansen have applied these types of ideas to the geometry of voiced speech. Their paper *Intrinsic Fourier analysis on the manifold of speech sounds* [36] shows an efficient algorithm for creating a Fourier basis on the manifold of possible speech signals whose elements correspond to the phonemic content of speech. Their paper gives an example of using this type of basis to create a spectrogram for the word “advantageous” that more efficiently identifies its phonemes than a conventional Fourier based spectrogram.

So far as we know, nobody has applied these mathematical developments to research aimed at improving hearing aids. Past results suggest that such methods are very likely to improve noise reduction. It is less certain, however, that such methods will improve intelligibility, but because these wavelets are closely matched to the structure of speech at the phonemic level, they offer some hope of success.

We have found in our research that even small distortions of the speech component of a noisy signal can render the denoised speech no more understandable than the original noisy signal. The only way to determine if the method improves intelligibility is to test it with hearing impaired listeners. In the end, it is only the individual user who can tell you if the hearing aid is helpful. A mathematical theory can tell you that according to certain mathematical assumptions, a particular threshold is optimal, or that according to a set of objective measures, one algorithm compares favorably with alternative methods; but only through subject testing can you find out if the method improves intelligibility. We hope that our research will prove helpful in providing some insight into some of the challenges facing the use of wavelet type bases in hearing aid research as well as offering some guidelines about conducting such research and testing the validity of its findings.

References

1. *Special Issue on diffusion maps and wavelets*, Appl. Comput. Harmon. Anal. **21**, (2006).
2. J. Berger, R. Coifman, M. Goldberg, *Removing noise from music using local trigonometric bases and wavelet packets*, Jour. Audio Eng. Soc., 808-818, (1994).

3. R. Bilger, J. Nuetzel, W. Rabinowitz and C. Rzczowski, *Standardization of a test of speech perception in noise*, J. Speech Hear. Res., **27**, 32-48 (1984).
4. S. Boll, *Suppression of acoustic noise in speech using spectral subtraction*, IEEE Trans. Acoustics, Speech, and Sig. Proc. **27**, 113-120, (1979).
5. S. Boll, *Speech enhancement in the 1980's: noise suppression with pattern matching*. In S. Furui and M. Sondhi, editors, *Advances in Speech Signal Processing*, Marcel Dekker, 309-326.
6. G. Box and G. Jenkins, *Time Series Analysis-Forecasting and Control*, Holden Day, San Francisco, Ca., (1970).
7. J.C. Bremer, R.R. Coifman, M. Maggioni, A.D. Szlam, *Diffusion wavelets packets*, Appl. Comput. Harmon. Anal. **21**, 95-112, (2006).
8. R. Brey, M. Robinette, D. Chabries and R.W. Christiansen, *Improvement in speech intelligibility in noise employing an adaptive filter with normal and hearing-impaired subjects*, Jour. Rehab. Res. and Dev. **24**, 75-86, (1987).
9. D. Byrne and H. Dillon, *The National Acoustics Laboratories' (NAL) new procedure for selecting the gain and frequency response of a hearing aid*, Ear and Hearing, vol. **7**, 257-265, 1986.
10. D. Chabries, R.W. Christiansen, R. Brey, M. Robinette, R. Harris, *Application of adaptive digital signal processing to speech enhancement for the hearing impaired*, Jour. Rehab. Res. and Dev., **24** 65-74, (1987).
11. H. Chipman, E. Kolaczky and R. McCulloch, *Signal de-noising using adaptive Bayesian wavelet shrinkage*, Proc. IEEE-SP Intl. Symp. Time-Freq, Time Scale Anal., 225-228, (1996).
12. R. Coifman and D. Donoho, *Translation-invariant de-noising*. In A. Antoniadis, editor, *Wavelets and Statistics*, Springer-Verlag.
13. R. Coifman, M. Maggioni, *Diffusion wavelets*, Appl. Comput. Harmon. Anal. **21**, 53-94, (2006).
14. R. Coifman and F. Majid, *Adaptive waveform analysis and denoising*. In Y. Meyer and S. Roques, editors, *Progress in Wavelet Analysis and Applications*, 63-76, (1993).
15. R. Coifman and N. Saito, *Local discriminant bases and their applications*, J. Math. Imag. Vision **5** (1995) 337-358.
16. R. Coifman and V. Wickerhauser, *Entropy-based algorithms for best basis selection*, IEEE Trans. Inf. Theory **38**, (1992), 713-738.
17. I. Daubechies, *Orthonormal bases of compactly supported wavelets*, Comm. on Pure and Appl. Math. **4**, (1988) 909-996.
18. D. Donoho, *Unconditional bases are optimal bases for data compression and for statistical estimation*, Appl. Comput. Harmonic Analysis **1**, (1993) 100-115.
19. D. Donoho and I. Johnstone, *Ideal spatial adaptation via wavelet shrinkage*, Biometrika **81**, 425-455, (1994).
20. D. Donoho, I. Johnstone and G. Kerkycharian and D. Picard *Wavelet Shrinkage: Asymptopia?* J. Royal Stat. Soc. Ser. B. **2** 301-337, (1995).
21. A.J. Duquesnoy and R. Plomp, *The effect of a hearing aid on the speech-reception threshold of hearing-impaired listeners in quiet and in noise*, Jour. Acoust. Soc. Amer., vol. **83**, pp. 2166-2173, 1983.
22. L.A. Drake, J.C. Rutledge and J. Cohen, *Wavelet Analysis in Recruitment of Loudness Compensation*, IEEE Transactions on Signal Processing, Dec.1993.
23. C.W. Dunnett, *A multiple comparison procedure for comparing several treatments to a control*, Jour. Amer. Stat. Assoc., vol. **50**, 1096-1121. 1955.
24. P. Duchnowski and P.M. Zurek, *Villchur revisited: another look at automatic gain control of simulation of recruiting hearing loss*, J. acoust. Soc. Amer. **98**, 6, pp. 3170-3181 (1995).
25. Y. Ephraim and D. Malah, *Speech enhancement using a minimum mean-square error short-time spectral amplitude estimator*, IEEE ATrans. Acoust. Speech Sig. Proc. **32**, (1984) 1109-1122.
26. Y. Ephraim, D. Malah and B. Juang, *On the applications of hidden Markov models for enhancing noisy speech*, IEEE ATrans. Acoust. Speech Sig. Proc. **32**, (1984) 1846-1856.
27. Y. Ephraim and H.L. Van Trees, *A signal subspace approach for speech enhancement*, Proc. ICASSP 1993, volume **2**, (1993) 355-358.

28. Y. Ephraim and H.L. Van Trees, *A signal subspace approach for speech enhancement*, IEEE Trans. Speech Audio Proc. **3**, (1995) 251-266.
29. Y. Ephraim and H.L. Van Trees, M. Nilsson and S. Soli, Enhancement of noisy speech for the hearing impaired using the signal subspace approach. In Proc. NIH Interdisciplinary Forum on Hearing Aid Research and Development (1996).
30. D. Graupe, J. Grosspietsch, and S. Basseas, *A single-microphone-based self-adaptive filter of noise from speech and its performance evaluation*. Jour. Rehab. Res. Dev. **24**, (1987) 119-126.
31. R. Gray, *On the asymptotic eigenvalue distribution of Toeplitz matrices*, IEEE Trans. Inf. Theory **18**, (1972) 725-730.
32. J.P. Gagne, *Excess masking among listeners with a sensorineural hearing loss*, Jour. Acoust. Soc. Amer., vol. **81**, pp. 2311-2321, 1988.
33. J. Greenberg and P. Zurek, *Evaluation of an adaptive beamforming method for hearing aids*, J. Acoust. Soc. Amer. **91**, 1662-1676, (1992).
34. C.S. Hallpike and J.D. Hood, *Observation upon the neurological mechanism of the loudness recruitment phenomenon*, Acta Oto-Laryng. **50**, pp. 472-486, (1959).
35. IEEE, *IEEE recommended practice for speech quality measurements*, IEEE Trans. Audio Electroacoustics, pp. 227-246, (1969).
36. A. Jansen and P. Niyogi, *Intrinsic Fourier analysis on the manifold of speech sounds*, Proc. International Conference on Acoustics, Speech, and Signal Processing, Toulouse, France, 2006.
37. M. Jansen, *Noise reduction by wavelet thresholding*, Springer Verlag, (2001).
38. J. Kates, *Speech enhancement based on a sinusoidal model*, J. Speech Hear. Res. **37**, 449-464, (1994).
39. M.C. Killion, *The K-Amp hearing aid: an attempt to present high fidelity for persons with impaired hearing*, Amer. J. Audiology, pp. 52-74, (1993)
40. S. Kochkin, *MarkeTrak III: Why 20 million in US do not use hearing aids for their hearing loss*, Hearing J. **46**, pp. 1-8, (1993).
41. S. Kochkin, *MarkeTrak III: 10-year Customer satisfaction trends in the US Hearing Instrument Market*, Hearing Rev. **9**, pp. 1-8, (2002).
42. S. Kochkin, *MarkeTrak VII: Customer satisfaction with hearing instruments in the digital age*, Hearing J. **58**, pp. 30-39, (2005).
43. H. Levitt, *A historical perspective on digital hearing aids: how digital technology has changed modern hearing aids*, Trends in Amplification **11**, pp. 7-24, (2007).
44. M. Lang, H. Guo, J. Odegard, J. Burrus and R. Wells, *Non-linear processing of a shift invariant DWT for noise reduction*, IEEE Sig. Proc. Letters **3**, 10-12, (1995).
45. H. Levitt, M. Bakke, J. Kates, A. Neuman, T. Schwander and M. Weiss, *Signal processing for hearing impairment*, Scand. Audiol. **Supplement 38**, 7-19, (1993).
46. J. Lim, editor, *Speech Enhancement*, Prentice Hall, aenglewood Cliffs, N.J. (1983).
47. J. Lim and A. Oppenheim, *All-pole modeling of degraded speech*, IEEE Trans. Acoust. Speech Sig. Proc. **26**, 197-209, (1978).
48. H. Levitt and S.B. Resnick, *Speech reception by the hearing impaired: methods of testing and the development of new tests*, Scand. Audiol. Supp. **6**, pp. 107-130, (1978).
49. J. Makhoul and R. McAulay, *Removal of noise from noise-degraded speech*, Technical report, National Academy of Sciences, National Academy Press, Washington, D.C. (1989).
50. S. Mallat, *A theory for multiresolution signal decomposition: the wavelet representation*, IEEE Trans. Pattern Anal. and Machine Intell. **11**, 674-693, (1989).
51. R. McAulay and M. Malpass, *Speech enhancement using a soft-decision noise suppression filter*, IEEE Trans. Acoust. Speech Sig. Proc. **28**, 137-145, (1980).
52. E.S. Martin and J.M. Pickett, *Sensironeural hearing loss and upward spread of masking*, Jour. Speech Hear. Res., vol. **13**, pp. 426-437, 1970.
53. J.C. Pesquet, H. Krim and H. Carfantan, *Time-invariant orthonormal wavelet representations*, IEEE Trans. Sig. Proc. **44**, (1996) 1964-1970.
54. P. Peterson, N. Durlach, W. Rabinowitz and P. Zurek, *Multimicrophone adaptive beamforming for interference reduction in hearing aids*, J. Rehab. Res. Dev. **24**, 102-110, (1987).
55. J. Porter and S. Boll, *Optimal estimators for spectral resoration of noisy speech*. In Proc. ICASSP 1984, volume **1**, page 18A.2.

56. S. Quackenbush, T. Barnwell and M. Clements, *Objective measures of Speech Quality*, Prentice Hall, Englewood Cliffs, N.J. (1993).
57. J.C. Rutledge, *Speech Enhancement for Hearing Aids*, in Time-Frequency and Wavelet Transforms in Biomedicine, Metin Akay, Editor, IEEE Press, 1997.
58. J. Rissanen, *Modeling by shortest data description*, Automatica **14**, 465-471, (1978).
59. T. Roos, P. Myllymaki and J. Rissanen, *MDL Denoising Revisited*, IEEE Transactions on Signal Processing, vol. **57**, No. 9, 3347-3360, (2009).
60. N. Saito, *Simultaneous noise suppression and signal compression using a library of orthonormal and signal compression using a library of orthonormal bases and the minimum description length criterion*. In E. Foufoula-Georgiou and P. Kumar, editors, Wavelets in Geophysics, academic Press (1994).
61. C. Sammeth and M. Ochs, *A review of current "noise reduction" hearing aids: rationale, assumptions and efficacy*, Ear and Hearing **12**, 116S-124S (1991).
62. T. Schwander and H. Levitt, *Effect of two-microphone noise reduction on speech recognition by normal-hearing listeners*, Jour. rehab. res. and Dev. **24**, (1987) 87-92.
63. R. Schwartz, M. Berouti and J. Makhoul, *Enhancement of speech corrupted by acoustic noise*, ICASSP-79 Proceedings, page 208 (1979).
64. D. Sinha and A. Tewfik, *Low bit rate transparent audio compression using adapted wavelets*, IEEE Trans. Sig. Proc. **41**, 3463-3479, (1993).
65. M. Smith and T. Barnwell, *Exact reconstruction techniques for tree-structured subband coders*, IEEE Trans. Acoust.Sig. Proc. **34**, 434-441, (1986).
66. C. Stein, *Estimation of the mean of a multivariate normal distribution*, Ann. Stat. **9**, 6, pp. 1135-1151, (1981).
67. R. Tyler and F. Kuk, *The effects of "noise suppression" hearing aids on sonsonant recognition in speech-babble and low-frequency noise*, Ear and hearing **10**, 243-249, (1989).
68. A.R. Thornton and M.J.M. Raffin, *Speech discrimination scores modeled as a binomial variable*, Jour. Speech Hear. Res., vol. **21**, pp. 507-518, 1978.
69. B. Vidakovic, *Non-linear wavelet shrinkage with Bayes rules and Bayes factors*, Discussion Paper 94-24, Duke University (1994).
70. J. Verschure and P.P.G. Van Benthem, *Effect of hearing aids on speech perception in noisy situations*, Audiology, vol. **31**, pp. 205-221, 1992.
71. W. Voiers, *Diagnostic acceptability measure for speech communication systems*. In Proc. ICASSP 1977, pages 204-207, (1977).
72. S. Watanabe, *Karhunen-Loève expansion and factor analysis: Theoretical remarks and applications*, Trans. 4th Prague Conf. Inform. Theory, Statist. Decision Functions, rand. Proc., 635-660, Prague Publishing House of the czechoslovak Academy of Sciences, (1967).
73. N. Whitmal, *Wavelet-Based Noise reduction for Speech Enhancement*, PhD thesis, Northwestern University, Evanston, IL (1997).
74. N. Whitmal, J. Rutledge and J. Cohen, *Reducing correlated noise in digital hearing aids*, IEEE Eng. Med. Biol. Mag. **15**, 88-96, (1996).
75. N. Whitmal, J. Rutledge and J. Cohen, *Reduction of autoregressive noise with shift-invariant wavelet packets*, Proc. IEEE-SP Symp. Time-Freq. Time-Scale Analysis, pp. 137-140, (1996).
76. G. Wornell (1990). *A Karhunen-Loève-like expansion for 1/f processes via wavelets*, IEEE Trans. Inf. Theory **36**, 859-861, (1990).
77. N. Whitmal and A. Vosoughi, *Recruitment of loudness effects of attenuative noise reduction algorithms*, J. Acoust. Soc. Amer. vol. **111**, Issue 5, p. 2380, (2002), (Conference Proceedings).

Index

- anomaly detection, 190
- autoregressive noise, 305, 311, 312
- autoregressive process, 308

- band-limited functions, 264, 272, 274, 275
- Besov spaces, 264, 265
- Bessel
 - norm, 141, 149
 - vector, 138–141, 143, 145–147
 - wavelet, 132, 143, 144
 - multiplier, 143, 144
- Bessel generator
 - multipliers, 131, 132, 134, 135, 138–141, 149
 - algebra of, 134
 - vectors, 131, 132, 134
- Bessel sequence, 132–134
 - generators, 136
 - strongly disjoint, 133
- Bezout identity, 118–121
- bracket product, 69

- Calderón reproducing formula, 261
- CMB polarization, 260, 273
- compact Riemannian manifold, 215
- composite dilation wavelet, 83–86, 90, 94, 95, 99, 104, 105
- composite scaling function, 85, 86, 88, 91, 92, 94, 96, 97, 100, 105
- compression, 162, 163, 168, 170, 190, 193, 194
- continuous wavelet transform, 279, 282–284
- cosmic microwave background radiation, 259, 260, 267, 268, 272–274
- coupled geometry, 163–168, 170, 173, 184, 187, 190, 191, 194, 196
 - of questionnaires and databases, 165
- crystallographic group, 83, 84, 87, 88, 90, 92, 97, 99–102, 104, 106
 - fundamental region of, 88
- crystallographic Haar wavelets, 102
- crystallographic Haar-type composite scaling function, 95
- cubature points, 264, 272

- data matrix, 161–165, 173, 176, 184, 186, 190–192, 196
- data-dependent dictionaries, 212
- Daubechies' condition, 261
- denoising, 1, 2, 9, 13, 128, 163, 191, 199, 200, 220, 300, 301, 305, 308, 310, 320, 327
- Density and Connectivity of Wavelet Frames, 135
- Diagnostic Acceptability Measure, 314
- diffusion distance, 229, 236, 237, 247
- diffusion maps, 229, 231, 236, 238, 245–248, 253, 254
- diffusion operator, 229, 236
- diffusion wavelet packets, 328
- dimension estimation, 199, 203
- dimensionality reduction, 227–229, 231, 250, 252–254
- discrete Fourier transform, 301
- discrete wavelet transform, 302
- distance measure, 227, 228, 230–232, 241, 243, 246, 247
- dual affinity, 167
- dual lattice, 67
- dyadic orthonormal wavelets, 3
- dyadic orthonormal wavelets, 3, 17–20

- Earth mover's distance, 227, 230, 241, 243
- earth mover's distance, 231, 253

- elliptic operators, 265, 266
- Empirical Orthogonal Function, 291
- entropy condition, 161, 165
- entropy function, 307
- expansive
 - automorphism, 36, 39, 40, 45, 46, 48
 - matrices, 19, 21
- Fourier Bessel wavelet multiplier, 144, 145
- frame bounds, 133, 136
- frame generator vector, 132–134
- frame wavelets, 7, 28
- full rank condition, 110
- full rank interpolatory schemes, 110
- Gaussian family, 273
- Gaussian function, 290, 292
- geometric harmonics multiscale extension, 238
- geometric wavelets, 199, 203, 211, 213, 215, 217
- global partition function, 289
- Hölder h type singularities, 287, 289
- Hölder exponent, 288, 289
- Haar filter, 79
- Haar MRA, 79, 80
- Haar-type composite dilation wavelet, 83, 84, 86, 89, 90
- Hausdorff distance, 231, 241, 249, 253
- hearing impaired, 300, 317, 325, 327, 328
- high pass filter, 88, 94–96, 98, 100, 104
- interpolatory subdivision scheme, 113
- Journé wavelet set, 3
- k -nearest neighbor, 230, 246
- Karhunen–Loève basis, 301
- Karhunen–Loève Transform, 229
- Laplace–Beltrami operator, 263, 267, 270
- Laplacian eigenfunctions, 162, 163, 170, 196
- Laplacian eigenmap, 227, 229, 231, 234–238, 241, 244, 245, 248–250, 252–254
- Laplacian eigenvectors, 161, 167, 170
- lattice, 66–68, 72
- law of large numbers, 274, 275
- Legendre transform, 289
- lip reading, 252
- local commutant, 138, 148
- locally compact Abelian groups, 3, 4, 17, 22, 35
- low pass filter, 71, 72, 76, 77, 79, 80, 88, 92–95, 97, 98
- Low-Dimensional Embeddings, 227, 231, 243
- Markov random walk, 229, 246
- masking, 299
- matrix refinement equation, 114
- matrix wavelets, 110
- Mexican hat function, 262
- Mexican needlets, 260, 265, 274, 275
- minimum description length, 308, 310
- missing value imputation, 163, 190
- Monge–Kantorovich optimal mass transfer, 242
- Monte-Carlo sampling, 220
- multichannel
 - multiresolution analysis, 110, 113, 115
 - wavelet, 109–111, 113, 115–119, 121–123, 125, 127–129
 - wavelet construction, 118
- multidimensional scaling, 227, 229, 231, 232
- multifractal signal, 287, 289
- multiresolution analysis, 3, 4, 6, 18, 23, 37, 65, 69, 70, 85, 86, 136
 - on the torus, 69
- multiscale dimension estimation, 204
- multiscale geometric analysis, 199, 212
- multiscale partition tree, 161, 163, 166
- nearly tight frame, 259–261, 263, 271
- neighborhood-mapping construction, 26
- noise reduction, 299, 301, 302, 305, 308, 313, 314, 316–318, 325, 328
- nonlinear manifold, 201, 213
- normalized graph Laplacians, 235
- Nyström extension, 238, 239
- orthonormal (respectively Riesz) basis
 - generator vector, 132
- p -adic
 - absolute value, 35
 - field, 20, 35
 - integers, 17, 35
 - rational numbers, 17, 35

- Parseval frame, 133, 144, 146, 148
 - set, 143, 144
 - vector, 142, 146
- Parseval frame wavelets, 28
- partition tree, 163, 166, 167, 170, 171, 173, 175, 183, 185, 191, 194
- partition-based correlation, 170, 171
- Poisson summation formula, 156, 265
- principal component analysis, 200, 203, 227, 229, 231–233
- projective unitary representation, 145–148
- pseudo-spectral representation, 281
- pseudo-translations, 37, 38
- pseudoinverse, 240

- quadrature filter bank, 302
- Quasi-Biennial Oscillation, 279, 280, 282–286
- Quincunx
 - dilation matrix, 66
 - matrix, 65
 - wavelets, 65, 67, 69, 71, 73, 75, 77, 79, 81

- recruitment of loudness, 300, 321, 327
- refinable functions, 69–71
- refinement mask, 114
- reproducing kernel Hilbert space, 230
- Revised Speech Perception in Noise Test, 302, 317–320
- Riemannian
 - manifold, 259, 260, 263
 - measure, 263
 - metric, 260, 267

- scaling function, 85, 86, 90, 92–97, 99, 100, 102–104, 106
- second order isotropy condition, 273
- Shannon filter, 77
- Shannon MRA, 77
- Shannon Sampling Theorem, 157
- shift-invariant spaces, 66, 68, 69
- shift-invariant wavelet-packet transform, 306, 310, 311, 313, 315, 318, 319
- signal-to-noise ratio, 299–303, 311, 317, 320
- signature construction, 244
- signatures of ensembles, 243
- Singular Value Decomposition, 200

- singularity spectrum, 287, 289, 292
- spectral theorem, 263
- spectrum of singularity, 287, 289
- spherical needlets, 259, 272, 274
- spin spherical harmonics, 270, 274
- Spin wavelets, 267
- spin-lowering operator, 270
- spin-raising operator, 270
- Stokes parameters, 268, 269
- subspace wavelet set, 31

- Tensor Haar-like bases, 161, 173, 185, 191
- tight frame, 133
 - generator, 133
- tiling of, 3, 21, 22, 24, 28, 30
- tiling set, 84, 87, 99, 100
- training ensembles, 227, 228, 230, 237, 241, 245, 248, 252

- unitary system, 131–135, 137–140, 142, 143, 145, 146, 148, 149
- unnormalized graph Laplacian, 235

- vector subdivision
 - scheme, 110
- vector subdivision
 - operator, 111
 - scheme, 109, 111–113, 115, 123

- wandering vector multipliers, 133, 134
- wandering vectors, 131
- wavelet frames, 132, 133, 135, 136
- wavelet sets, 2–4, 17–27, 29–35, 37, 39, 41, 43–47, 49, 51, 53, 55, 57, 59, 62, 63
- wavelet-packet transform, 302, 320
- wavelets
 - p -adic, 36
 - Littlewood-Paley, 23
 - on the torus, 73
 - Shannon, 17, 23, 24, 32, 48
- Wedding sets, 33
- Weil-Brezin mapping, 155

- Zak transform, 151, 153, 155–157

Applied and Numerical Harmonic Analysis

- J.M. Cooper: *Introduction to Partial Differential Equations with MATLAB* (ISBN 978-0-8176-3967-9)
- C.E. D'Attellis and E.M. Fernández-Berdaguer: *Wavelet Theory and Harmonic Analysis in Applied Sciences* (ISBN 978-0-8176-3953-2)
- H.G. Feichtinger and T. Strohmer: *Gabor Analysis and Algorithms* (ISBN 978-0-8176-3959-4)
- T.M. Peters, J.H.T. Bates, G.B. Pike, P. Munger, and J.C. Williams: *The Fourier Transform in Biomedical Engineering* (ISBN 978-0-8176-3941-9)
- A.I. Saichev and W.A. Woyczyński: *Distributions in the Physical and Engineering Sciences* (ISBN 978-0-8176-3924-2)
- R. Tolimieri and M. An: *Time-Frequency Representations* (ISBN 978-0-8176-3918-1)
- G.T. Herman: *Geometry of Digital Spaces* (ISBN 978-0-8176-3897-9)
- A. Procházka, J. Uhlíř, P.J.W. Rayner, and N.G. Kingsbury: *Signal Analysis and Prediction* (ISBN 978-0-8176-4042-2)
- J. Ramanathan: *Methods of Applied Fourier Analysis* (ISBN 978-0-8176-3963-1)
- A. Teolis: *Computational Signal Processing with Wavelets* (ISBN 978-0-8176-3909-9)
- W.O. Bray and C.V. Stanojević: *Analysis of Divergence* (ISBN 978-0-8176-4058-3)
- G.T. Herman and A. Kuba: *Discrete Tomography* (ISBN 978-0-8176-4101-6)
- J.J. Benedetto and P.J.S.G. Ferreira: *Modern Sampling Theory* (ISBN 978-0-8176-4023-1)
- A. Abbate, C.M. DeCusatis, and P.K. Das: *Wavelets and Subbands* (ISBN 978-0-8176-4136-8)
- L. Debnath: *Wavelet Transforms and Time-Frequency Signal Analysis* (ISBN 978-0-8176-4104-7)
- K. Gröchenig: *Foundations of Time-Frequency Analysis* (ISBN 978-0-8176-4022-4)
- D.F. Walnut: *An Introduction to Wavelet Analysis* (ISBN 978-0-8176-3962-4)
- O. Bratteli and P. Jorgensen: *Wavelets through a Looking Glass* (ISBN 978-0-8176-4280-8)
- H.G. Feichtinger and T. Strohmer: *Advances in Gabor Analysis* (ISBN 978-0-8176-4239-6)
- O. Christensen: *An Introduction to Frames and Riesz Bases* (ISBN 978-0-8176-4295-2)
- L. Debnath: *Wavelets and Signal Processing* (ISBN 978-0-8176-4235-8)
- J. Davis: *Methods of Applied Mathematics with a MATLAB Overview* (ISBN 978-0-8176-4331-7)
- G. Bi and Y. Zeng: *Transforms and Fast Algorithms for Signal Analysis and Representations* (ISBN 978-0-8176-4279-2)
- J.J. Benedetto and A. Zayed: *Sampling, Wavelets, and Tomography* (ISBN 978-0-8176-4304-1)
- E. Prestini: *The Evolution of Applied Harmonic Analysis* (ISBN 978-0-8176-4125-2)
- O. Christensen and K.L. Christensen: *Approximation Theory* (ISBN 978-0-8176-3600-5)
- L. Brandolini, L. Colzani, A. Iosevich, and G. Travaglini: *Fourier Analysis and Convexity* (ISBN 978-0-8176-3263-2)
- W. Freeden and V. Michel: *Multiscale Potential Theory* (ISBN 978-0-8176-4105-4)
- O. Calin and D.-C. Chang: *Geometric Mechanics on Riemannian Manifolds* (ISBN 978-0-8176-4354-6)

Applied and Numerical Harmonic Analysis (Cont'd)

J.A. Hogan and J.D. Lakey: *Time-Frequency and Time-Scale Methods* (ISBN 978-0-8176-4276-1)

C. Heil: *Harmonic Analysis and Applications* (ISBN 978-0-8176-3778-1)

K. Borre, D.M. Akos, N. Bertelsen, P. Rinder, and S.H. Jensen: *A Software-Defined GPS and Galileo Receiver* (ISBN 978-0-8176-4390-4)

T. Qian, V. Mang I, and Y. Xu: *Wavelet Analysis and Applications* (ISBN 978-3-7643-7777-9)

G.T. Herman and A. Kuba: *Advances in Discrete Tomography and Its Applications* (ISBN 978-0-8176-3614-2)

M.C. Fu, R.A. Jarrow, J.-Y. J. Yen, and R.J. Elliott: *Advances in Mathematical Finance* (ISBN 978-0-8176-4544-1)

O. Christensen: *Frames and Bases* (ISBN 978-0-8176-4677-6)

P.E.T. Jorgensen, K.D. Merrill, and J.A. Packer: *Representations, Wavelets, and Frames* (ISBN 978-0-8176-4682-0)

M. An, A.K. Brodzik, and R. Tolimieri: *Ideal Sequence Design in Time-Frequency Space* (ISBN 978-0-8176-4737-7)

S.G. Krantz: *Explorations in Harmonic Analysis* (ISBN 978-0-8176-4668-4)

G.S. Chirikjian: *Stochastic Models, Information Theory, and Lie Groups, Volume 1* (ISBN 978-0-8176-4802-2)

C. Cabrelli and J.L. Torrea: *Recent Developments in Real and Harmonic Analysis* (ISBN 978-0-8176-4531-1)

B. Luong: *Fourier Analysis on Finite Abelian Groups* (ISBN 978-0-8176-4915-9)

M.V. Wickerhauser: *Mathematics for Multimedia* (ISBN 978-0-8176-4879-4)

P. Massopust and B. Forster: *Four Short Courses on Harmonic Analysis* (ISBN 978-0-8176-4890-9)

O. Christensen: *Functions, Spaces, and Expansions* (ISBN 978-0-8176-4979-1)

J. Barral and S. Seuret: *Recent Developments in Fractals and Related Fields* (ISBN 978-0-8176-4887-9)

O. Calin, D.-C. Chang, K. Furutani, and C. Iwasaki: *Heat Kernels for Elliptic and Sub-elliptic Operators* (ISBN 978-0-8176-4994-4)

C. Heil: *A Basis Theory Primer* (ISBN 978-0-8176-4686-8)

J. Klauder: *A Modern Approach to Functional Integration* (ISBN 978-0-8176-4790-2)

J. Cohen and A.I. Zayed: *Wavelets and Multiscale Analysis* (ISBN 978-0-8176-8094-7)