Compressive Imaging: Structure, Sampling, Learning

Accurate, robust and fast image reconstruction is a critical task in many scientific, industrial and medical applications. Over the last decade, image reconstruction has been revolutionized by the rise of compressive imaging. It has fundamentally changed the way modern image reconstruction is performed. This in-depth treatment of the subject commences with a practical introduction to compressive imaging, supplemented with examples and downloadable code, intended for readers without extensive background in the subject. Next, it introduces core topics in compressive imaging – including compressed sensing, wavelets and optimization – in a concise yet rigorous way, before providing a detailed treatment of the mathematics of compressive imaging. The final part is devoted to the most recent trends in compressive imaging: deep learning and neural networks. This highly timely component provides, for the first time, a readable overview of these nascent topics. With an eye to the next decade of imaging research, and using both empirical and mathematical insights, it examines the potential benefits and the pitfalls of these latest approaches.

Ben Adcock is Associate Professor of Mathematics at Simon Fraser University. He received the CAIMS/PIMS Early Career Award (2017), an Alfred P. Sloan Research Fellowship (2015) and a Leslie Fox Prize in Numerical Analysis (2011). He has published 15 conference proceedings, two book chapters and over 50 peer-reviewed journal articles. His work has been published in outlets such as *SIAM Review* and *Proceedings of the National Academy of Sciences*, and featured on the cover of *SIAM News*.

Anders C. Hansen is Reader in Mathematics at the University of Cambridge and Professor of Mathematics at the University of Oslo. He received the Leverhulme Prize in Mathematics and Statistics (2017), the 2018 IMA Prize in Mathematics and Applications and the Whitehead Prize (2019). He has had papers published in outlets such as the *Journal of the American Mathematical Society* and *Proceedings of the National Academy of Sciences*, and featured on the cover of *Physical Review Letters* and *SIAM News*.

Compressive Imaging: Structure, Sampling, Learning

BEN ADCOCK
Simon Fraser University

ANDERS C. HANSEN
University of Cambridge

with contributions by
VEGARD ANTUN

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

79 Anson Road, #06–04/06, Singapore 079906

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of
education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781108421614
DOI: 10.1017/9781108377447

First published 2021

Printed in Singapore by Markono Print Media Pte Ltd

A catalogue record for this publication is available from the British Library.

ISBN 978-1-108-42161-4 Hardback

For

Tina

Gyda, Anna and Christian

Contents

	Preface		*page* xiii
1	**Introduction**		1
	1.1	Imaging and Inverse Problems	2
	1.2	What is Compressive Imaging?	3
	1.3	Terminology	5
	1.4	Imaging Modalities	5
	1.5	Conventional Compressed Sensing	7
	1.6	Imaging with Compressed Sensing	12
	1.7	Neural Networks and Deep Learning for Compressive Imaging	17
	1.8	Overview and Highlights	22
	1.9	Disclaimers	23
	1.10	Reading this Book	25
Part I	**The Essentials of Compressive Imaging**		27
	Summary of Part I		29
2	**Images, Transforms and Sampling**		30
	2.1	Images	30
	2.2	Sampling with the Fourier Transform	33
	2.3	Sampling with the Radon Transform	36
	2.4	Binary Sampling with the Walsh Transform	40
3	**A Short Guide to Compressive Imaging**		47
	3.1	The Six Stages of the Compressive Imaging Pipeline	47
	3.2	First Example: MRI	51
	3.3	Second Example: X-ray CT	58
	3.4	Third Example: Optical Imaging	62
	3.5	Recovery of Image Sequences	66
4	**Techniques for Enhancing Performance**		75
	4.1	Sampling Strategies	75
	4.2	How to Design a Sampling Strategy	80
	4.3	Resolution Enhancement	82

4.4 Discretization, Model Mismatch and the Inverse Crime 86
4.5 Sparsifying Transforms 88
4.6 Iterative Reweighting 90
4.7 Learning: Incorporating Training Data 93
4.8 An All-Round Model-Based Compressive Imaging Strategy 97

Part II Compressed Sensing, Optimization and Wavelets 101
Summary of Part II 103

5 An Introduction to Conventional Compressed Sensing 105
5.1 Projections 105
5.2 Sparsity and Compressibility 106
5.3 Measurements and Measurement Matrices 108
5.4 ℓ^1-Minimization 112
5.5 Recovery Guarantees in Compressed Sensing 114
5.6 Techniques for Uniform Recovery 115
5.7 Techniques for Nonuniform Recovery 123
5.8 Oracle Estimators 125

6 The LASSO and its Cousins 129
6.1 Definitions 129
6.2 Summary 130
6.3 Uniform Recovery Guarantees 133
6.4 Nonuniform Recovery Guarantees 135

7 Optimization for Compressed Sensing 142
7.1 Minimizers, Not Minimum Values 142
7.2 Flavours of Optimization for Compressive Imaging 145
7.3 Gradient Descent 147
7.4 Forward–Backward Splitting 148
7.5 The Primal–Dual Iteration 151
7.6 Nesterov's Method and NESTA for QCBP 157

8 Analysis of Optimization Algorithms 166
8.1 The rNSP and Inexactness 166
8.2 Compressed Sensing Analysis of the Primal–Dual Iteration 169
8.3 Compressed Sensing Analysis of NESTA 171
8.4 Computability and Complexity of Finding Minimizers 174
8.5 Impossibility of Computing Minimizers from Inexact Input 177
8.6 Complexity Theory for Compressed Sensing Problems 178

9 Wavelets 188
9.1 Introduction 188
9.2 Multiresolution Analysis 192

9.3 Wavelet Construction from an MRA 193
9.4 Wavelet Design 198
9.5 Compactly Supported Wavelets 200
9.6 Daubechies Wavelets 202
9.7 Wavelets on Intervals 207
9.8 Higher Dimensions 210
9.9 Summary and Orderings 213
9.10 Discrete Wavelet Computations 214
9.11 Beyond Wavelets 219

10 A Taste of Wavelet Approximation Theory 222
10.1 Linear and Nonlinear Approximation 222
10.2 Linear and Nonlinear Wavelet Approximation Rates 224
10.3 Discussion 225
10.4 Proofs of the Approximation Results 228
10.5 Higher Dimensions 234

Part III Compressed Sensing with Local Structure 237
Summary of Part III 239

11 From Global to Local 241
11.1 The Fourier-Wavelets Problem 241
11.2 Structured Sparsity and the Flip Test 244
11.3 Local Sparsity in Levels 250
11.4 Sampling Operators 252
11.5 Notions of Coherence and Recovery Guarantees 259
11.6 The One-Dimensional Discrete Fourier–Haar Wavelet Problem 262

12 Local Structure and Nonuniform Recovery 267
12.1 Weighted ℓ^1-Minimization 268
12.2 Local Recovery Guarantee 269
12.3 The Sparse Model 270
12.4 The Sparse in Levels Model 271
12.5 The Fourier–Haar Wavelet Problem 276
12.6 Comparison with Oracle Estimators 280
12.7 Coherences of the Discrete Fourier–Haar Matrix 281
12.8 Proof of Theorem 12.4 284
12.9 Improvements for Random Signal Models 297

13 Local Structure and Uniform Recovery 305
13.1 Decoders and Recovery Guarantees 305
13.2 Restricted Isometry and Robust Null Space Properties 307
13.3 Stable and Accurate Recovery via the rNSPL and G-RIPL 308
13.4 Measurement Conditions for Uniform Recovery 314
13.5 Proof of Theorem 13.12 320

14 Infinite-Dimensional Compressed Sensing 334
14.1 Motivations 334
14.2 When Vectors Become Functions 336
14.3 Bandwidth 337
14.4 Generalized Sampling 338
14.5 Compressed Sensing in Infinite Dimensions 341
14.6 Recovery Guarantees for Weighted QCBP 343
14.7 Feasibility and the SR-LASSO 344

Part IV Compressed Sensing for Imaging 349
Summary of Part IV 351

15 Sampling Strategies for Compressive Imaging 353
15.1 Overview 353
15.2 The DS, DIS and DAS Schemes 358
15.3 Summary of the Measurement Conditions 360
15.4 When and Why Do Certain Patterns Work Well? 367

16 Recovery Guarantees for Wavelet-Based Compressive Imaging 373
16.1 Decoders and Recovery Guarantees 373
16.2 What does this Mean? Near-Optimal Wavelet Approximation 374
16.3 Main Recovery Guarantees 378
16.4 Proofs of the Recovery Guarantees 381

17 Total Variation Minimization 403
17.1 Definitions 403
17.2 Two Curious Experiments 405
17.3 Recovery Guarantees for Fourier Sampling 406
17.4 Proofs 411
17.5 Recovery Guarantees for Haar-Incoherent Measurements 419
17.6 Structure-Dependent Sampling 422

Part V From Compressed Sensing to Deep Learning 427
Summary of Part V 429

18 Neural Networks and Deep Learning 431
18.1 Supervised Machine Learning 431
18.2 Neural Networks 432
18.3 The Universal Approximation Theorem 435
18.4 Architecture Design and Extensions 436
18.5 Training a Neural Network 439
18.6 The Success of Deep Learning 445
18.7 Instabilities in Deep Learning 446
18.8 Why do Instabilities Occur? 452

19 Deep Learning for Compressive Imaging 458
19.1 Why Go Beyond Compressed Sensing 458
19.2 Deep Learning for Inverse Problems 459
19.3 Experimental Setup 465
19.4 Testing for Instabilities in Neural Networks for Compressive Imaging 466

20 Accuracy and Stability of Deep Learning for Compressive Imaging 470
20.1 Optimal Maps 470
20.2 Examining Instabilities in Deep Learning for Compressive Imaging 475
20.3 The Stability versus Performance Tradeoff 484
20.4 The Tradeoff for Compressed Sensing and Deep Learning 487
20.5 Can Instabilities be Remedied? 493

21 Stable and Accurate Neural Networks for Compressive Imaging 501
21.1 Algorithms for Computing Neural Networks 501
21.2 Unravelled ISTA 503
21.3 Unravelling the Primal–Dual Iteration 505
21.4 Stable and Accurate Neural Networks for Compressed Sensing 509
21.5 Stable and Accurate Neural Networks for Compressive Imaging 513

Epilogue 521

Appendices 525
Summary of the Appendices 527

Appendix A Linear Algebra 528
A.1 Norms 528
A.2 Orthogonality and Orthonormal Bases 529
A.3 Matrices 530
A.4 Matrix Norms 531
A.5 Further Properties of the Matrix ℓ^2-Norm 532
A.6 The Singular Value Decomposition 533
A.7 Least Squares and the Pseudoinverse 533
A.8 Circulant Matrices and Convolution 535
A.9 Kronecker products of matrices 535

Appendix B Functional Analysis 537
B.1 Metric Spaces 537
B.2 Banach and Hilbert Spaces 538
B.3 The ℓ^p and L^p Spaces 538
B.4 Operators between Normed Spaces 539
B.5 Orthogonality in Hilbert Spaces 540
B.6 Orthonormal Bases of Hilbert Spaces 540

Cambridge University Press
978-1-108-42161-4 — Compressive Imaging: Structure, Sampling, Learning
Ben Adcock , Anders C. Hansen
Frontmatter
More Information

Appendix C Probability 542
C.1 Definition and Basic Properties 542
C.2 Random Variables 543
C.3 Joint Distributions and Independence 544
C.4 Expectation and Variance 544

Appendix D Convex Analysis and Convex Optimization 546
D.1 Convex Sets and Convex Functions 546
D.2 Subdifferentials 547
D.3 Convex Optimization 548
D.4 The Convex Conjugate 549
D.5 The Proximal Map 551

Appendix E Fourier Transforms and Series 553
E.1 The Fourier Transform 553
E.2 The Fourier Basis 554
E.3 The Discrete Fourier Transform 555

Appendix F Properties of Walsh Functions and the Walsh Transform 556
F.1 Sign Changes 556
F.2 Binary Addition, Translation and Scaling 556
F.3 Orthonormality and Relation to Haar Wavelets 558
F.4 The Paley and Sequency Orderings 560
F.5 Properties of the Hadamard Matrix 561
F.6 The Fast Walsh–Hadamard Transform 562

Notation 563
Abbreviations 567
References 568
Index 596

Preface

The objective of *compressive imaging* is to develop algorithms for image reconstruction that exploit the low-dimensional structure inherent to natural images to achieve higher-quality reconstructions from fewer measurements. This structure has been used in *image compression* since the 1990s. But it was not until the 2000s, which saw the advent of the related field of *compressed sensing* – introduced in the work of Candès, Romberg & Tao and Donoho – that it began to be exploited in the context of image reconstruction. Nowadays, compressive imaging is a large and vibrant subject, spanning mathematics, computer science, engineering and physics. It has fundamentally altered how images are reconstructed in a variety of real-world settings. Classical linear reconstruction techniques have in many cases been replaced by sophisticated nonlinear reconstruction procedures based on convex, or sometimes even nonconvex, optimization problems. Practical applications include Magnetic Resonance Imaging (MRI), X-ray Computed Tomography (X-ray CT), electron or fluorescence microscopy, seismic imaging and various optical imaging modalities, to name but a few. The field continues to evolve at a rapid rate, with the most recent trend being the introduction of tools from machine learning, such as deep learning, as means to achieve even further performance gains.

Objectives of this Book

This book is about compressive imaging and its mathematical underpinnings. It is aimed at graduate students, postdoctoral fellows and faculty in mathematics, computer science, engineering and physics who want to learn about modern image reconstruction techniques. Its goal is to span the gap between theory and practice, giving the reader both an overview of the main themes of compressive imaging and an in-depth mathematical analysis. A consistent theme of the book is the insight such mathematical analysis brings, both in designing methods in the first place and then enhancing their practical performance.

The book consists of 22 chapters, plus appendices containing various prerequisite materials. It is divided into five parts. Part I is a practical guide to compressive imaging, supported by many numerical examples and downloadable code. It is intended for readers without extensive background in the subject. Part II systematically introduces the main mathematical tools of compressive imaging, including conventional compressed sensing, convex optimization and wavelets. Parts III and IV are devoted to compressed sensing theory and its application to image reconstruction, respectively. Finally, Part V considers the most recent trends in compressive imaging, namely, deep learning and neural networks. It provides an overview of this nascent topic and includes both mathematical and empirical insights.

This book contains many numerical examples. A companion software library, *CIlib*, has been developed by Vegard Antun (University of Oslo). It contains a broad set of functions and tools for experimenting with various compressive imaging techniques. It is available at

`https://github.com/vegarant/cilib`

or through the book's website:

`www.compressiveimagingbook.com`

This library also includes code for reproducing most of the figures in the book.

Acknowledgements

Many people have contributed substantially to this book. First and foremost, we owe a great debt of gratitude to Vegard Antun. He has produced all the figures in the book and developed the aforementioned software library. He also provided vital input into many parts of this project. Suffice to say, without Vegard this would be a far worse book!

We also gratefully acknowledge the contributions of our (current and former) students and postdocs who have assisted with this project. They have proofread large parts of the manuscript, providing useful suggestions and pointing out both typos and more serious errors. They have also shown great patience during the final stages of its preparation as other projects suffered inevitable delays. They are (in alphabetical order): Edvard Aksnes, Anyi Bao, Alex Bastounis, Randall Bergman, Simone Brugiapaglia, Juan M. Cardenas, Il-Yong Chun, Matthew Colbrook, Nick Dexter, Einar Gabbassov, Milana Gataric, Nina Gottschling, Alex Jones, Matthew King-Roskamp, Chen Li, Zhen Ning David Liu, Mathias Lohne, Kristian Monsen Haug, Sebastian Moraga, Max Neyra-Nesterenko, Clarice Poon, Francesco Renna, Mohsen Seifi, Yi Sui, Laura Thesing and Qinghong Xu.

Many colleagues have also contributed to this book, through collaborating with us, discussing and answering our many questions, and proofreading parts of the manuscript. This long list includes (in alphabetical order): Simon Arridge, Richard Baraniuk, Stephen Becker, Claire Boyer, Robert Calderbank, Emmanuel Candès, Antonin Chambolle, Felipe Cucker, Ingrid Daubechies, Mike Davies, Ron DeVore, David Donoho, Yonina Eldar, Jalal Fadili, Alhussein Fawzi, Hamza Fawzi, Omar Fawzi, Jeffrey Fessler, Simon Foucart, Michael Friedlander, Pascal Frossard, Anne Gelb, Rémi Gribonval, Karlheinz Gröchenig, David Gross, Nick Higham, Bamdad Hosseini, Arieh Iserles, Mark Iwen, Laurent Jacques, Jakob Sauer Jørgensen, Felix Krahmer, Richard Kueng, Gitta Kutyniok, Bradley Lucier, Jackie Ma, Dustin Mixon, Mohsen Moosavi-Dezfooli, Amirafshar Moshtaghpour, Adrian Nachman, Arkadi Nemirovski, Nilima Nigam, Ozan Öktem, Hooman Owhadi, Gabriel Peyré, Yaniv Plan, Rodrigo Platte, Bogdan Roman, Justin Romberg, Øyvind Ryan, Rayan Saab, Carola Schönlieb, Thomas Strohmer, Andrew Stuart, Joel

Tropp, Michael Unser, Mihaela van der Schaar, Martin Vetterli, Verner Vlačić, Felix Voigtlaender, Michael Wakin, Rachel Ward, Clayton Webster, Pierre Weiss and Ozgur Yilmaz. We extend our thanks to all of them and to anyone whose name we have accidentally left off this list.

Several images were provided to the authors for use in the numerical experiments. We thank Elizabeth Sawchuk (and Bear) for the 'dog' image, Vegard Antun for the 'klubbe' and 'kopp' images, Andy Ellison for the 'pomegranate' image and GE Electric Healthcare for the 'brain' image. See Fig. 2.2. We also thank Dr. Cynthia McCollough, the Mayo Clinic, the American Association of Physicists in Medicine and the National Institute of Biomedical Imaging and Bioengineering for allowing the use of their data in the experiments shown in Figs 20.4 and 20.10. Table 18.1 and Figs 18.5 and 18.8 are from [359]. We thank the authors, Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi and Pascal Frossard, for allowing us to reproduce them. The deep learning experiments in Chapters 19 and 20 were first produced in [36]. We thank the authors of [253, 278, 427, 512, 527] for sharing their code, data and network weights, and thereby allowing us to reproduce their networks and experiments. We also thank Kristian Monsen Haug for his assistance with Fig. 21.1.

Both authors would like to acknowledge their respective institutions for their support: Purdue University, Simon Fraser University, University of Cambridge and University of Oslo. BA also wishes to thank the Isaac Newton Institute for several productive visits. BA acknowledges support from an Alfred P. Sloan Research Fellowship, NSERC (the Natural Sciences and Engineering Research Council of Canada) through grant 611675, NSF (the National Science Foundation) through grant DMS-1318894, PIMS (the Pacific Institute for the Mathematical Sciences) through the Collaborative Research Group 'High-dimensional Data Analysis' and SFU's Big Data Initiative through the 'Next Big Question Fund'. ACH acknowledges support from a Royal Society University Research Fellowship, EPSRC (the Engineering and Physical Sciences Research Council) through grant EP/L003457/1 and the 2017 Leverhulme Prize. We also acknowledge Cambridge University Press and, in particular, our editor David Tranah, who has been an invaluable source of guidance while writing this manuscript. Lastly, we thank our respective families for their continued love and support throughout this project.

Burnaby, BC, Canada	Ben Adcock
Cambridge, UK	Anders C. Hansen
August 2020	