Monographs in Computer Science

Editors

David Gries
Fred B. Schneider

Monographs in Computer Science

Abadi and Cardelli, **A Theory of Objects**

Alagic, **Rational Database Technology**

Bauer and Wössner, **Algorithmic Language and Program Development**

Benosman and Kang [editors], **Panoramic Vision: Sensors, Theory, and Applications**

Bhanu, Lin, and Krawiec, **Evolutionary Synthesis of Pattern Recognition Systems**

Bronstein, Bronstein, and Kimmel, **Numerical Geometry of Non-Rigid Shapes**

Broy and Stolen, **Specification and Development of Interactive Systems**

Brzozowzki and Seger, **Asynchronous Circuits**

Burgin, **Super-Recursive Algorithms**

Cantone, Omodeo, Policriti, and Schwartz, **Set Theory for Computing: From Decision Procedures to Declarative Programming with Sets**

Downey and Fellows: **Parameterized Complexity**

Feijen, Gasteren, Gries, and Misra [editors], **Beauty is our Business: A Birthday Salute to Edsger W. Dijkstra**

Feijen and Gasteren, **On a Method of Multiprogramming**

Gries, **The Science of Programming**

Gries and Schneider, **A Logical Approach to Discrete Math**

Grune and Jacobs, **Parsing Technique: A Practical Guide**

Hehner, **A Practical Theory of Programming**

Herbert and Spärck Jones [editors], **Computer Systems: Theory, Technology, and Applications**

Heydon, Levin, Mann, and Yu, **Software Configuration Management Using Vesta**

Kozen, **The Design and Analysis of Algorithms**

Mclver and Morgan, **Abstraction, Refinement, and Proof for Probabilistic Systems**

Mclver and Morgan [editors], **Programming Methodology**

Mishra, **Algorithmic Algebra**

Misra, **A Discipline of Multiprogramming: Programming Theory for Distributed Applications**

Partsch, **Specification and Transformation of Programs: A Formal Approach to Software Development**

Paton [editor], **Active Runes in Database Systems**

Poernomo, Crossley, and Wirsing, **Adapting Proofs-as-Programs: The Curry-Howard Protocol**

Preparata and Shamos, **Computational Geometry: An Introduction**

Randell [editor], **The Origins of Digital Computers: Selected Papers**

Reps and Teitelbaum, **The Synthesizer Generator Reference Manual**

Selig, **Geometric Fundamentals of Robotics**

Sasha and Zhu, **High Performance Discovery in Time Series: Techniques and Case Studies**

Tonella and Potrich, **Reverse Engineering of Object Oriented Code**

Vitányi, **Algorithmic Statistics**

Alexander M. Bronstein
Michael M. Bronstein
Ron Kimmel

Numerical Geometry of Non-Rigid Shapes

With 10 Color Figures

Springer

Alexander M. Bronstein
Technion-Israel Institute of Technology
Haifa, Israel
bron@cs.technion.ac.il

Michael M. Bronstein
Technion-Israel Institute of Technology
Haifa, Israel
mbron@cs.technion.ac.il

Ron Kimmel
Technion-Israel Institute of Technology
Haifa, Israel
ron@cs.technion.ac.il

ISSN: 0172-603x
ISBN: 978-0-387-73300-5 e-ISBN: 978-0-387-73301-2
DOI 10.1007/978-0-387-73301-2

Library of Congress Control Number: 2008934481

Mathematics Subject Classification (2000): 53AØ5, 52C25, 49M37

© 2008 Springer Science+Business Media, LLC
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 developedis 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

springer.com

The mediocre teacher tells.
The good teacher explains.
The great teacher inspires.

W. A. Ward

Dedicated to our teachers.

Foreword

Once upon a time, a child imagined a fierce boa constrictor that swallowed an elephant, giving rise to a most peculiar shape. The child made a drawing of it and showed it to grown-ups around him, expecting them to be appalled by the scene. Alas, nobody saw in his drawing more than a hat. The child then proceeded to draw an explanatory drawing showing the elephant inside the snake's expansible stomach...

Non-rigid objects are all around us, from snakes to octopuses, from ropes to the pages of this book, from the surface of the sea to the pudding in our plates, and we have no particular problems in dealing with them in our daily lives. Yet the mathematical tools we have for their description and analysis are few, and only relatively recently researchers in graphics and computer vision have started paying due attention to them. This book offers a rare opportunity to encounter the fascinating world of the flexible, elastic, plastic, and amorphous form and shape through the looking glass of their mathematical representation and numerical treatment by our miraculous and powerful computing machines.

I hope that after reading this book, you, too, just like the Little Prince, will start to see in Saint-Exupery's childhood drawing the elephant inside the boa constrictor, rather than the obvious, boring and rigid hat.

Alfred M. Bruckstein
Ollendorff Professor of Science
Technion, Haifa, Israel

About the Authors

Ron Kimmel is a researcher in the areas of computer vision, image processing, and numerical geometry. He received his Ph.D. from the Department of Electrical Engineering, Technion – Israel Institute of Technology in 1995 and spent the post-doctoral years (1995–1998) at the University of California in Berkeley. In 2003–2004, he was a visiting professor the Stanford University. Currently, he is a professor in the Department of Computer Science, Technion.

Ron has authored more than one hundred twenty articles in scientific journals and conferences and the book *Numerical Geometry of Images*. He is a member of the editorial board of *International Journal of Computer Vision* and of *IEEE Transactions on Image Processing* and was the co-chair of the *International Conference on Scale-Space Methods in Computer Vision* in 2001 and 2005 and of the *SIAM Conference on Imaging* in 2006. Prof. Kimmel is the recipient of the Hershel Rich Technion Innovation Award (twice), the Taub Prize for Excellence in Research, the Gensler Counter-Terrorism Prize, and the Alon Fellowship. He was a consultant to HP Labs and to Net2Wireless/Jigami research and currently serves on the advisory board of MediGuide.

Alexander and Michael Bronstein met Ron in 2002, taking his class on *Numerical Geometry of Images* when still undergraduate students. According to Michael's recollection, the first phrase Prof. Kimmel referred to him was "usually I throw such students out of the class" when his cellular phone accidentally rang during the lecture. Shortly afterwards, Alex and Michael became Ron's graduate students, working together on analysis of non-rigid objects, computer vision, and image processing problems. They each received a Ph.D. from the Technion in 2007.

The Bronstein twins are the alumni of the Technion Excellence Program and the Academy of Achievement. They shared numerous awards, including the Technion Humanities and Arts Department prize, Kasher prize, Thomas Schwartz award, Hershel Rich Technion Innovation award, Gensler Counter-Terrorism Prize, and the Best Paper award at the *Copper Mountain Conference on Multigrid Methods*. In 2003, their research on three-dimensional face recognition was a hit item in CNN news and other international media.

Michael and Alex were awarded the Adams Fellowship by the Israel Academy of Science in 2005 and 2006, respectively.

Besides this book, Alex, Michael, and Ron co-authored more than a dozen journal and conference papers, dedicated mostly to the analysis of three-dimensional non-rigid shapes, and co-founded Novafora Inc. Among colleagues and friends, the harmonic trio is frequently referred to as "B^2K."

Preface

Looking around, we notice that our world is full of objects that, due to their physical properties, are non-rigid and therefore can be deformed and bent. Non-rigid shapes appear at all scales in Nature – from the human body, its organs and tissues, to tiny bacteria and microscopic protein molecules. Being so ubiquitous, such shapes are often encountered in pattern recognition and computer vision applications. The richness of the possible deformations of non-rigid shapes appears to be a nightmare for a pattern recognition researcher, who faces a vast number of degrees of freedom when trying to analyze them. For this reason, explicit analysis of non-rigid objects has been avoided for a long period in computer vision, and as it often happens in applied sciences, research has focused on simplified problems that are easier to solve.

Today, there is a gradually penetrating comprehension that in many applications the necessity to model and understand non-rigid objects is unavoidable. At the same time, recent research results have shown that problems related to non-rigid objects are not necessarily untractable. We decided to write this book because we believed that a critical mass of research has accumulated, making the field of non-rigid shape analysis sufficiently profound on one hand and having significant open research questions on the other.

In some sense, the book can be considered as a sequel to *Numerical Geometry of Images*, which focused mainly on geometric methods in image processing and analysis. In *Numerical Geometry of Non-rigid Shapes*, as the title suggests, our main theme is two- and three-dimensional non-rigid objects. We invite the reader to join us for a fascinating journey to the non-rigid world, a rapidly developing field at the crossroad of computer vision, pattern recognition, and geometry, where the last word has not yet been spoken.

Intended use

This book was initially written as lecture notes for a one-semester monographic graduate course taught at the Technion – Israel Institute of Technology.

The course was based on a corpus of recent research results, and its leitmotif was intrinsic geometric invariants and embedding methods. However, because the field of non-rigid object analysis is multi-disciplinary, bringing together theoretical and numerical geometry, optimization, graph theory, machine learning, and computer graphics, it is almost impossible to expect students to have all the necessary background. Therefore, we had to extend and elaborate the text, including a gradual introduction of the material required for understanding the mathematical machinery we use throughout the book. The reader is assumed to have basic knowledge of calculus and algebra and can acquire the required background through reading the introductory chapters.

The book is intended as a graduate-level textbook for engineers, computer scientists, and applied mathematicians. For teachers, it can be the main reference for a monographic course on non-rigid shape analysis or a supplementary material for various courses in computer vision and pattern recognition, geometry processing, computational and numerical geometry, and computer graphics. For students, the book can be both course material and a self-study reference. For mature experts and specialists in the field, the book offers frontline methods and most recent results, as well as less traditional points of view on old problems and approaches. In the book and supporting online material, the reader can find numerical recipes, ready to use codes and references to public domain and commercial software.

Features

When working on the book, we set ourselves to three main goals. First, we tried to make the book as self-contained as possible. Given the breadth of the material covered, it is obvious that in-depth study of all the fields could easily spread over ten other books. One of the most difficult tasks was to cherry-pick only the relevant material and present it consistently. Second, we were convinced that any material, however complicated, could be simply presented. We therefore tried to maintain simplicity of explanation throughout the book, often resorting to illustrative descriptions, drawing inspiration in realms ranging from fairy tales to cartography. We believe this made the reading even of the driest technical material somewhat more entertaining. Finally, a special emphasis was made on practical applications, as our largest audience is engineers and applied scientists who would like to see things work at the end of the day. We tried to keep the right balance between the simplicity of explanation and the amount of details necessary for implementation of the discussed numerical algorithms. For those willing to dive deeper into additional details, we provide references to related books and research papers. In order to increase the practical value of the book, we also provide references to public domain and commercial software.

Each chapter includes references to implementation of some of the discussed algorithms and a list of suggested literature for those interested in

wider and deeper understanding. We believe this will be especially appreciated by readers using the book for self-education. In order to make the discussion as focused as possible, we omit certain details, concentrating them into notes at the end of each chapter. Proofs of some results are formulated as problems, on which the reader may test his or her understanding of the material. Solutions of selected problems appear at the end of the book; other problems are left as an exercise to the reader, which may help in using the book as supplementary class material. The level of problems varies significantly, from exercises involving basic calculus to open research questions. To leverage the background and terminology and facilitate the reading, we provide a glossary with the definitions of the most frequently used terms. At the end of the book, the reader can find a subject and an author index.

Synopsis

The first chapters are dedicated to the mathematical background necessary to create a common language and terminology. In many books, the theoretical introduction is often apparently unrelated to the problems discussed afterwards. We tried to avoid this impression, targeting our examples and illustrations at the problems we solve, thereby motivating the background chapters.

Our model of non-rigid world stands on three pillars: metric geometry, discrete geometry, and numerical optimization. We begin our mathematical introduction (Chapter 2) with geometry. Geometry of surfaces is typically presented either from the point of view of topology or by resorting to the heavy apparatus of calculus and differential geometry. This creates an apparent gap: whereas topology is too "crude" and usually does not satisfy our needs in studying surfaces, differential geometry is often cumbersome and poorly accessible by the average reader. We preferred to follow a somewhat less orthodox path and present the *metric* point of view, using notions as basic as distance and length. This, following Dmitry Burago's expression, brings geometry back "down-to-earth" where it traditionally, and literally, began. In our introduction to geometry, we explore the difference between *intrinsic* and *extrinsic* geometry, the notion of *isometry*, and invariant description of shapes. Chapters 3 and 4 present the geometric foundations through the glasses of a computer scientist, who has to convert the continuous geometric world into discrete objects tractable by a computer. Much attention in Chapter 4 is dedicated to *fast marching*, a class of efficient numerical algorithms for geodesic distance measurement. The third pillar, numerical optimization, is presented in Chapter 5.

Starting from Chapter 6, applications begin to appear. To ease the entrance into the non-rigid world, we first study a simple problem of rigid object analysis, which has a smaller number of degrees of freedom. Our emphasis in this chapter is on *iterative closest point* algorithms, a class of numerical

methods for rigid object comparison and matching. In Chapters 7, 8, and 9, we discuss ways of creating deformation-invariant representations of non-rigid shapes. The main focus is on *multidimensional scaling* methods, first developed in psychology in the 1960s for multidimensional data analysis, and more recently adapted to non-rigid surface matching. Chapter 10 is devoted to problems of shape similarity, which is addressed by constructing an ideal deformation-invariant distance. We encounter the Gromov-Hausdorff distance, introduced in the early 1980s by Mikhail Gromov and first used in pattern recognition applications by Facundo Mémoli and Guillermo Sapiro in 2005, and study its theoretical properties and approaches for approximate computation. In Chapter 11, we discuss *partial similarity*, the problem of comparing shapes having similar parts. Chapter 12 deals with the problem of correspondence between non-rigid shapes. This problem is presented following the same line of shape similarity. In Chapter 13, we study the problem of face recognition, one of the most challenging computer vision applications. Face recognition appears as a playground to demonstrate and summarize the tools discussed in the book. Finally, Chapter 14 gives a short retrospective and concludes the book.

How to use this book

The book is largely structured in such a way that new material builds up progressively and thus it can be read continuously. Some sections not essential for the understanding of the entire material or those containing challenging mathematics are denoted by $\star$. These sections can be skipped without sacrificing future understanding.

Readers new to the field and those using the book for self-education are recommended to start with the introductory chapters. Experts in the field may be interested in focused topics offering new insight and state-of-the-art methods. Those include parallel fast marching (Chapter 4), multigrid and vector extrapolation accelerated multidimensional scaling (Chapter 7), generalized multidimensional scaling (Chapter 9), Gromov-Hausdorff distance (Chapter 10), Pareto similarity (Chapter 11), and expression-invariant face recognition (Chapter 13).

Teachers may use the book for a monographic course on analysis and synthesis of non-rigid shapes, using the flow of the chapters. Detailed study of examples, some technical descriptions, and solutions to selected problems (marked with ✓) can be a basis for tutorials. Problems without solutions are a potential source for home assignments. Alternatively, Chapters 4, 9, and 12 may be used as supplementary material for computer graphics related courses and Chapters 7, 8, 10, and 11 for enrichment of computer vision and pattern recognition curricula.

Acknowledgment

The publication of this book would not have been possible without the help, advice, and support of many people to whom we are deeply indebted. Particularly, we would like to mention Ronen Basri, Alexander Brook, Alfred Bruckstein, Dmitry Burago, Laurent Cohen, Mathieu Desbrun, David Donoho, Yohai Dvir, Ilya Eckstein, Michael Elad, Gene Golub, Hugues Hoppe, Martin Kilian, Nahum Kiryati, Vladimir Lin, Nathan Linial, Facundo Mémoli, Gérard Medioni, Davide Migliore, Gabriel Peyré, Boaz Porat, Helmut Pottmann, Dan Raviv, Guy Rosman, Guillermo Sapiro, Michael Saunders, Peter Schröder, James Sethian, Avraham Sidi, Nir Sochen, Alla Sheffer, Alon Spira, Pierluigi Taddei, Ofir Weber, Irad Yavneh, and Michael Zibulevsky, whose valuable comments helped us shape the manuscript at different stages of its life. Special thanks to Constantine Hadavas, Alex Hubris, Idan Shatz, Ayellet Tal, and Malene Thyssen for the permission to reproduce some of the illustrations used in the book. We are also grateful to the Springer team, Vaishali Damle and Ann Kostant, for helping us create the book you are currently reading.

Feedback and support

Because non-rigid shape analysis is an active and fast evolving research field, *Numerical Geometry of Non-rigid Shapes* is intended as a living text. We will do our best to keep the book up to date. Updates, new materials, data, and software will be published on the website `tosca.cs.technion.ac.il/book`. Comments, corrections, suggestions, and interesting solutions to problems would always be highly appreciated.

Haifa and Santa Clara,
April 2008

Alexander M. Bronstein
Michael M. Bronstein
Ron Kimmel

Contents

Foreword VII

About the Authors IX

Preface XI

1 Introduction 1
1.1 Similarity of non-rigid shapes 3
1.2 Correspondence problems 6
1.3 A landscape of problems 7
Notes 9

2 A Taste of Geometry 11
2.1 Basic terms in metric geometry and topology 11
2.2 Isometries 13
2.3 Length spaces 17
2.4 Manifolds 20
2.5 Embedded surfaces 21
2.6 Curvature and the second fundamental form 26
2.7 Intrinsic view on geometry of surfaces 29
2.8 Bending and rigidity 31
2.9 Intrinsic invariants 34
Suggested reading 37
Problems 37
Notes 39

3 Discrete Geometry 41
3.1 Point clouds and sampling 41
3.2 Farthest point sampling 43
3.3 Voronoi tessellation 46
3.4 Centroidal Voronoi sampling and the Lloyd-Max algorithm 48
3.5 Connectivity 52
3.6 Delaunay tessellation 53
3.7 Triangular meshes 54
3.8⋆ Local feature size and curvature-dependent sampling 57
3.9⋆ Approximation quality 61
Suggested reading 63
Software 63
Problems 63
Notes 64

4 Shortest Paths and Fast Marching Methods 67
4.1 The shortest path problem 67
4.2 Dijkstra's shortest path algorithm 69
4.3 Fast marching methods 71
4.4 Fast marching on parametric surfaces 81
4.5 Marching even faster 83
4.6 Parallel distance computation 85
4.7⋆ Minimal geodesics 87
Suggested reading 89
Software 90
Problems 90
Notes 91

5 Numerical Optimization 93
5.1 Local versus global optimization 93
5.2 Optimality conditions 94
5.3 Unconstrained optimization algorithms 97
5.4 The quest for a descent direction 100
5.5 Preconditioning 104
5.6 Let Newton be! 105
5.7 Truncated Newton 106
5.8 Quasi-Newton algorithms 107
5.9 Non-convex optimization 108
5.10 Constrained optimization 110
5.11 Penalty and barrier methods 112
5.12⋆ Augmented Lagrangian method 114
Suggested reading 116
Software 116
Problems 116
Notes 118

6 In the Rigid Kingdom . . . 119
6.1 Moments of joy, moments of sorrow . . . 120
6.2 Iterative closest point algorithms . . . 125
6.3 Enter numerical optimization . . . 128
6.4 Rigid correspondence . . . 131
Suggested reading . . . 133
Software . . . 133
Problems . . . 133
Notes . . . 134

7 Multidimensional Scaling . . . 137
7.1 Isometric embedding problem . . . 138
7.2 Multidimensional scaling . . . 142
7.3 SMACOF algorithm . . . 143
7.4* Second-order methods . . . 146
7.5 Variations on the stress theme . . . 148
7.6 Multiresolution methods . . . 153
7.7* Multigrid MDS . . . 156
7.8* Vector extrapolation . . . 160
7.9 A trouble with topology . . . 164
Suggested reading . . . 165
Software . . . 166
Problems . . . 167
Notes . . . 167

8 Spectral Embedding . . . 169
8.1 Classic MDS . . . 170
8.2 Local methods . . . 173
8.3 The Laplace-Beltrami operator . . . 176
8.4 To hear the shape of the drum . . . 178
8.5* Discrete Laplace-Beltrami operator . . . 180
Suggested reading . . . 184
Software . . . 184
Problems . . . 184
Notes . . . 185

9 Non-Euclidean Embedding . . . 187
9.1 Spherical embedding . . . 187
9.2 Generalized multidimensional scaling . . . 192
9.3 Representation issues . . . 194
9.4 Geodesic distance computation . . . 197
9.5 Minimization of the generalized stress . . . 198
9.6 Multiresolution encore . . . 202
Suggested reading . . . 203
Software . . . 203

Problems 204
Notes 204

10 Isometry-Invariant Similarity 205
10.1 Equivalence, similarity, and distance 205
10.2 Embedding distance 207
10.3 Gromov-Hausdorff distance 208
10.4 Intrinsic symmetry 211
Suggested reading 214
Problems 214

11 Partial Similarity 217
11.1 Recognition by parts 218
11.2 Paretian approach to partial similarity 221
11.3 Scalar partial similarity 224
11.4 Fuzzy approximation 226
11.5 Extrinsic partial similarity 229
11.6 Intrinsic partial similarity 230
11.7 Not only size matters 232
Suggested reading 236
Problems 236
Notes 238

12 Non-rigid Correspondence and Calculus of Shapes 239
12.1 Intrinsic parameterization 240
12.2 An image processing approach 241
12.3 Minimum distortion correspondence 244
12.4 Texture mapping and transfer 246
12.5 Morphing 249
12.6⋆Guaranteed self-intersection free morph 254
12.7 Calculus of shapes 255
Suggested reading 258
Software 258
Problems 259
Notes 259

13 Three-dimensional Face Recognition 261
13.1 Some terminology 263
13.2 A retrospective 264
13.3 Isometric model of facial expressions 268
13.4 Expression-invariant face recognition 269
13.5 Comparison of photometric properties 273
Suggested reading 275
Notes 275

14 Epilogue 277

Solutions of Selected Problems 279

Software 293

Notation 297

Acronyms 299

Glossary 301

References 307

Subject Index 327

Author Index 335