Preface

> It is better to solve the right problem the wrong way than to solve the wrong problem the right way.
>
> R. W. Hamming

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 front-line 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, Gerard 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