Preface

Work on this book started more than 15 years ago, when I began a revision of a textbook from 1974 on numerical methods. That book devoted only about 90 pages to matrix computations compared to the more than 700 pages of the present book. This difference reflects not only a change in ambition, but also an increase in size and importance of the subject. A stunning growth in hardware performance has allowed more sophisticated mathematical models to be employed in sciences and engineering. In most of these applications, solution of systems of linear equations and/or eigenvalue problems lies at the core. Increased problem sizes and changes in computer architecture have also made the development of new methods and new implementations of old ones necessary.

Although there is a link between matrix computations and linear algebra as taught in departments of mathematics, there are also several important differences. Matrices are used to represent many different objects such as networks and images, besides linear transformations. Concepts such as ill-conditioning, norms, and orthogonality, which do not extend to arbitrary fields, are central to matrix computations. This is the reason for not using "linear algebra" in the title of the book.

This book attempts to give a comprehensible and up-to-date treatment of methods and algorithms in matrix computations. Both direct and iterative methods for linear systems and eigenvalue problems are covered. This unified approach has several advantages. Much of the theory and methods used to solve linear systems and eigenvalue problems are closely intertwined—it suffices to think of matrix factorizations and Krylov subspaces.

It is inevitable that personal interests would to some extent influence the selection of topics. This is most obvious in Chap. 2, which gives an unusually broad coverage of least squares methods. Several nonstandard topics are treated, e.g., tensor problems, partial least squares, and least angle regression. Methods for solving discrete inverse problems are also treated. Nonlinear least squares problems such as exponential fitting, nonlinear orthogonal regression, and logistic regression are covered. Parts of this chapter were originally written for a never published revised edition of my 1996 monograph entitled *Numerical Methods for Least Squares.*

The book is suitable for use in a two-semester course on matrix computations at advanced undergraduate or graduate level. The first semester could cover direct methods for linear systems and least squares using Chaps. 1 and 2; the second semester eigenvalue problems and iterative methods using Chaps. 3 and 4. But other combinations are possible. As prerequisite, a basic knowledge of analysis and linear algebra and some experience in programming and floating point computations will suffice. The text can also serve as a reference for graduates in engineering and as a basis for further research work. Problems and computer exercises are included after each section. It is highly recommended that a modern interactive system such as Matlab be available for working out these assignments. It should be stressed that the Matlab programs included in the text are mainly for illustration. They work, but are toy programs and not in any way close to production codes.

To keep the book within reasonable bounds, complete proofs are not given for all theorems. For the pursuit of particular topics in more detail, the book contains a large comprehensive and up-to-date bibliography of historical and review papers, as well as recent research papers. Care has been taken to include references to the original research papers since these are often rewarding to read. More than 50 short biographical notes on mathematicians who have made significant contributions to matrix computations are given as footnotes in the text.

When working on this book I soon realized that I was trying to hit a moving target. Having rewritten one chapter I invariably found that some other chapter now needed to be revised. Therefore, many draft versions have existed. At various stages of this process several colleagues, including Bo Einarsson and Tommy Elfving, read parts of early drafts and made many constructive comments. I am also greatly indebted to Michele Benzi, Nick Higham, and David Watkins, as well as several anonymous reviewers, whose suggestions led to major improvements in later versions of the text. I am indebted to Wlodek Proskurowski for his continuous encouragement and for using early versions of the book for courses at USC, Los Angeles.

Michael Saunders somehow found time to proofread the last draft and painstakingly corrected my faulty English language and other lapses. Without his help I would not have been able to get the book in shape. Finally, I thank Lynn Brandon, my editor at Springer, for her helpful and professional support during the publication process.

The book was written in Emacs and typeset in LATEX, the references were prepared in BibTEX, and the index with MakeIndex. Matlab was used for working out examples and generating figures. Using these great tools is always a joy. The biographical notes are based on the biographies compiled at the School of Mathematics and Statistics, University of St Andrews, Scotland (www-history.mcs.st-andrews.ac.uk).

The book is dedicated to the memory of George E. Forsythe and Gene H. Golub, who were my gracious and inspiring hosts during my postdoctoral stay at Stanford University in 1968. Their generosity with their time and ideas made

this stay into a decisive experience for my future life and work. Finally, I thank my wife Eva for her forbearance and understanding during all the time I spent on writing this book.

Linköping, November 2013   Åke Björck