Preface

This book is intended to be used as a text for beginning graduate-level (or even senior-level) students in mathematics, the sciences, engineering, computer science, or computational science who wish to acquire a working knowledge of computational matrix analysis or numerical linear algebra in a one-quarter or one-semester course. By matrix analysis I mean the use of linear algebra and matrix theory together with their intrinsic interaction with systems of linear constant coefficient differential or difference equations. Some of the fundamentals of this theory are described in the book *Matrix Analysis for Scientists and Engineers* by the same author. The adjective "computational" (or perhaps "qualitative" would be better) in the title of this book means the numerical aspects of matrix analysis. In that sense, the present text serves a follow-on role, either as a formal course or for self-study. Students should learn enough in this textbook that they can be aware of and at least ask the right questions in their subsequent use of numerical linear algebra in day-to-day computing or research.

The choice of topics covered in computational matrix analysis or numerical linear algebra is, for the most part, fairly standard. After a brief review of notation and some of the topics in conventional matrix analysis, an introduction to finite (IEEE) arithmetic is given. This is what it's all about: how to compute the things we want under the very real constraint of doing it in finite arithmetic. This is followed by a discussion, mostly by examples, of conditioning, stability, and rounding analysis, as well as an introduction to some mathematical software topics related to numerical linear algebra. This is followed by a thorough introduction to Gaussian elimination (LU factorization). Solution of nonsingular linear systems via Gaussian elimination is discussed along with condition estimation techniques, including statistical condition estimation. A chapter on linear least squares covers orthogonal reduction and the QR factorization. The next chapter discusses eigenvalue/eigenvector problems, some perturbation theory, and algorithms for their solution via the QR algorithm. Variants of the QR algorithm for singular value decomposition (SVD) and generalized eigenvalue problems follow which implement still more matrix factorizations. Applications of some of the algorithms discussed are then given in a final chapter. These include methods for the solution of matrix Lyapunov, Sylvester, and Riccati equations. In all cases, enough of the features of the principal algorithms are discussed so that the student can delve into the myriad details in the extant literature.

Prerequisites for using this text are quite modest, essentially just calculus and definitely some previous exposure to matrices and linear algebra at the level of *Matrix Analysis for Scientists and Engineers*. In particular, topics such as pseudoinverses and the SVD should be very familiar. These powerful and versatile tools are exploited to provide a unifying foundation upon which many of the topics are based. Moreover, because tools such as the SVD are not generally amenable to "hand computation," this approach necessarily presupposes the availability of appropriate mathematical software on a digital computer. For this, I highly recommend MATLAB although other (cheaper) software (such as Octave or Scilab) is also excellent. It is also important to recognize that the pervasive impact of MATLAB or lookalikes on the mathematics, science, and engineering communities has been as much on *communication* as on *computation*. Generally, the only way to describe an algorithm effectively is through its reliable implementation as mathematical software. How much easier the introduction of numerical linear algebra tools such as the SVD or the generalized eigenvalue problem would have been in the control and systems community in the mid-1970s had such software been available!

The material in this textbook can be thought of as "matrix analysis for grown-ups" for at least two reasons. First, "real-life" problems seldom yield to simple closed-form formulas or solutions. They must generally be solved computationally and it is important to know which types of algorithms can be relied upon and which cannot. Some of the key algorithms of classical numerical linear algebra, in particular, form the foundation upon which rests virtually all of modern computational science and engineering. A second motivation is that it provides many of the essential tools for what might be called "qualitative mathematics." For example, in an elementary linear algebra course, a set of vectors is either linearly independent or it is not. This is an absolutely basic concept. But in most engineering or scientific contexts we want to know more than that. If a set of vectors is linearly independent, how "nearly dependent" are the vectors? If they are linearly dependent, are there "best" linearly independent subsets? Such questions, many of which are related to the computation of rank, turn out to be much more difficult problems and frequently involve research-level questions when set in the context of having to be computed in the finite-precision, finite-range floating-point arithmetic environment of most modern computing platforms.

Some of the applications of computational matrix analysis mentioned briefly in this book are based on the modern state-space approach to dynamical systems. Instructors are encouraged to supplement the book with specific application examples from their own particular subject areas. State-space methods are now standard in much of modern engineering where, for example, control systems with large numbers of interacting inputs, outputs, and states often give rise to models of very high order that must be analyzed, simulated, and evaluated. The "language" in which such models are most conveniently described involves vectors and matrices. It is thus crucial that students have a working knowledge of the vocabulary and grammar of this "language." The tools of matrix analysis and numerical linear algebra are also applied on a daily basis to problems in biology, chemistry, econometrics, physics, statistics, and a wide variety of other fields, and thus the text can serve a rather diverse audience. Mastery of the material in this text should prepare the student to read and understand the use of matrix analysis and numerical linear algebra throughout mathematics, science, and engineering.

I have taught this material for many years, mostly at UCSB, twice at UC Davis, and several times at UCLA. The book is based on the course Math 270B, which is taught at UCLA and is designed to cover dense numerical linear algebra. A follow-on course, Math 270C, covers many of the same topics but concentrates on the case of large, sparse matrices or matrices with exploitable structure. Clearly, there is much more to the subject than this short book can possibly do justice to. It provides merely a taste of what is available, and I hope this introduction to the subject will make the reader at least a more discerning consumer of the algorithms and software that are out there. A more thorough knowledge of numerical linear algebra requires, at a minimum, a more complete treatment of rounding analysis than is provided here and most definitely requires more than a one-quarter course. I recommend that interested students take, either formally or informally, any of a number of advanced courses. After studying from this textbook, they will at least be prepared to do so.

Finally, it is a pleasure to acknowledge a debt of gratitude to G.W. (Pete) Stewart. His 1973 textbook [38] was my first introduction to numerical linear algebra, and I have borrowed considerably from it. That text, although nearing forty years old, has stood the test of time well.

— AJL, November 2011