

# Preface to the 1996 Edition

---

The following are natural goals for an aspiring textbook writer of a book like this one:

1. It should be attractive to first year graduate students from a variety of engineering and scientific disciplines.
2. The text should be self-contained, assuming only a good undergraduate background in linear algebra.
3. The students should learn the mathematical basis of the field, as well as how to build or find good numerical software.
4. Students should acquire practical knowledge for solving real problems efficiently. In particular, they should know what the state-of-the-art techniques are in each area, or when to look for them and where to find them.
5. It should all fit in one semester, since that is what most students have available for this subject.

The fifth goal is perhaps the hardest to manage. The first edition of these notes was 215 pages, which did fit into one ambitious semester. This edition has more than doubled in length, which is certainly too much material for even a heroic semester. The new material reflects the growth in research in the field in the last few years, including my own. It also reflects requests from colleagues for sections on certain topics that were treated lightly or not at all in the first edition. Notable additions include

- a class homepage with Matlab source code for examples and homework problems in the text, and pointers to other on-line software and textbooks;
- more pointers in the text to software for all problems, including a summary of available software for solving sparse linear systems using direct methods;
- a new chapter on Krylov subspace methods for eigenvalue problems;
- a section on domain decomposition, including both overlapping and nonoverlapping methods;

- sections on “relative perturbation theory” and corresponding high-accuracy algorithms for eigenproblems, like Jacobi and qd;
- more detailed performance comparisons of competing least squares and symmetric eigenvalue algorithms; and
- new homework problems, including many contributed by Zhaojun Bai.

A reasonable one-semester curriculum could consist of the following chapters and sections:

- all of Chapter 1;
- Chapter 2, excluding sections 2.2.1, 2.4.3, 2.5, 2.6.3, and 2.6.4;
- Chapter 3, excluding section 3.5;
- Chapter 4, up to and including section 4.4.5;
- Chapter 5, excluding sections 5.2.1, 5.3.5, 5.4 and 5.5; and
- Chapter 6, excluding sections 6.3.3, 6.5.5, 6.5.6, 6.6.6, 6.7.2, 6.7.3, 6.7.4, 6.8, 6.9.2, and 6.10.

Homework problems are marked Easy, Medium or Hard, according to their difficulty. Problems involving significant amounts of programming are marked “programming”.

Many people have helped contribute to this text, notably Zhaojun Bai, Alan Edelman, Velvel Kahan, Richard Lehoucq, Beresford Parlett, and many anonymous referees, all of whom made detailed comments on various parts of the text. Table 2.2 is taken from the PhD thesis of my student Xiaoye Li. Alan Edelman at MIT and Martin Gutknecht at ETH Zurich provided hospitable surroundings at their institutions while this final edition was being prepared. Many students at Courant, Berkeley, Kentucky and MIT have listened to and helped debug this material over the years, and deserve thanks. Finally, Kathy Yelick has contributed scientific comments, latex consulting, and moral support over more years than either of us expected this project to take.

James Demmel  
MIT  
September, 1996