

nonlinear eigenvalue problems, notions of Fredholm operators, periodic solutions of Hamiltonian systems, saddle points of nonquadratic functionals are among the important topics not discussed.

The mathematical community owes a debt of thanks to Fucik and Kufner, the two Czech authors of this book. They have produced a readable account of important contemporary topics in nonlinear analysis. These days, so much important research of our best people is dribbled out of them, piecemeal, in the form of imperfectly developed journal articles and conference proceedings. Let us hope that in the near future, other highly talented mathematicians of nonlinear science will be afforded the opportunity and leisure to share with us their finest conceptions in the form of systematically developed books, accessible to a wide mathematically educated audience.

M. S. BERGER

BULLETIN (New Series) OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 4, Number 3, May 1981
© 1981 American Mathematical Society
0002-9904/81/0000-0214/\$02.50

The symmetric eigenvalue problem, by Beresford N. Parlett, Prentice-Hall Series in Computational Mathematics, Prentice-Hall, Englewood Cliffs, N. J., 1980, xix + 348 pp., \$ 25.00 cloth.

The thesis of this review may be summarized in three propositions. First, numerical analysis is a science with mathematical, empirical, and engineering components. Second, a conventional mathematical education does not equip one to deal with the last two components. Third, the book under review is a good place for a mature mathematician to get an appreciation of all three aspects of the subject.

At the outset I would like to correct a possible misapprehension. For most of this essay, I am going to focus on the nonmathematical aspects of numerical analysis. This does not mean that I wish to minimize the role of mathematics in numerical analysis; on the contrary, it is hard to overstate its importance. But the pure mathematician coming to the field for the first time will find much that is strange, and I hope this review will provide a brief guide to this extra-mathematical territory.

The mathematical component of numerical analysis scarcely needs arguing. The subject derives its analytic tools from many branches of mathematics. Its journals usually present results in the form of theorems, the coin by which mathematical productivity is currently measured. Nor are these theorems more trivial or less rigorously established than those of other branches of mathematics. Finally, most numerical analysis courses are listed in mathematics departments, perhaps jointly with a computer science department.

The empirical component of numerical analysis derives from the fact that numerical analysis is a branch of applied mathematics, and its results are therefore subject to outside verification. In general, an applied mathematician must look on a piece of experimental apparatus with a mixture of hope and trepidation, since it can confirm or deny his researches with unarguable

finality. The numerical analyst's particular *bête noir* is, paradoxically, the machine that is responsible for the present flowering of the subject—the modern digital computer. The most involved mathematical investigation of an algorithm is subject to instant refutation by a high school student playing around with a home computer. However, since numerical analysts are not generally in the habit of proving false theorems, the refutation usually takes the form of a demonstration that the analysis does not say much that is useful about the algorithm.

The notion of a meaningful theorem is so important in numerical analysis that I will discuss it at some length. The point of this excursion is not that there are good and bad theorems—something that is true of all branches of mathematics—but that the standards by which a theorem is judged come from outside mathematics. For purposes of illustration, I shall state a hypothetical theorem and then consider how it can fail to be useful.¹

Suppose we are given a function $\phi: R^n \rightarrow R^n$ which is defined in terms of some numerical data and wish to determine a fixed point of ϕ by an iteration of the form $x_{k+1} = \phi(x_k)$. Consider the following theorem.

THEOREM. *For any value of the input data the following is true.*

1. *There is a unique x^* such that $\phi(x^*) = x^*$.*
2. *For any x_0 the sequence defined by $x_{k+1} = \phi(x_k)$ ($k = 0, 1, 2, \dots$) converges to x^* .*
3. *There is a constant $K > 0$ such that*

$$\lim_{k \rightarrow \infty} \frac{\|x_{k+1} - x^*\|}{\|x_k - x^*\|^2} = K. \quad (1)$$

In many respects this would appear to be an ideal theorem. The first statement insures the existence of the object we wish to compute. The second guarantees that the iteration is global; i.e., it converges to x^* from any starting point. The third statement asserts the Q -quadratic convergence of the algorithm. Informally, it may be taken as saying that the components of x_{k+1} will ultimately have twice as many correct digits as those of x_k .

Let us now see what this theorem fails to say. In the first place, it says nothing about the dependence of x^* on the input data. If x^* varies wildly with small variations of the input data, then the computed solution need have no discernable relation to the true one, except in those rare cases where the data are known exactly. Note that we are not talking about ill-posed problems, which are essentially discontinuous; rather we are concerned with nominally continuous problems whose behavior is so bad that they are practically intractable. The need to detect such *ill-conditioned* problems is the reason why perturbation theory plays a large role in modern numerical analysis.

The global convergence vouched for by the second statement will be of little use if the iteration takes too long to get near x^* . In fact, many of the

¹ The theorem is hypothetical in name only; similar theorems suffering from many of the defects listed below have actually been published.

simplest iteration schemes for solving numerical problems are, with unimportant exceptions, globally convergent (e.g. Bernoulli's method for finding a zero of a polynomial); but many of these are infrequently used because the convergence can be too slow.

A major difficulty with the third statement is that it does not give a bound on K . If K is large, then the region within which the iteration exhibits the behavior characteristic of quadratic convergence will be small, perhaps smaller than the accuracy required of the computed solution.

Another problem is that the wrong choice of norm in (1) can make the quadratic convergence meaningless. Typically, one wants an equable norm that does not weight individual components unduly, say one of the Hölder norms ($p = 1, 2, \infty$ are favorites here, as everywhere else). However, it frequently happens that the natural norm in which to establish a theorem about a numerical process is unbalanced. In finite-dimensional spaces the temptation is to appeal to the equivalence of all norms, which, for example, would assure us that Q -quadratic convergence is invariant under change of norm. Unfortunately, the constant K is not, and the passage from an unreasonable norm to a reasonable one is likely to make K unreasonably large.

Finally, the theorem says nothing about how the iteration behaves in the presence of rounding error. The limitations of finite precision computation have been the undoing of many promising algorithms. An example, which is treated in the book under review, is the Lanczos algorithm. This extraordinarily powerful method for approximating eigenvalues of a symmetric matrix waited in limbo for thirty years until a way was found to control the effects of rounding errors on it.

Up to now we have been concerned with how a theorem can fail to reveal how bad an algorithm is. There is a converse problem; many algorithms are better than existing analyses prove them to be. This may be because the complexities of the algorithm render a complete, rigorous analysis impossible, in which case one must be content with suggestive theorems, often about special cases. There are also some well-known algorithms, such as Gaussian elimination with partial pivoting, which cannot be shown to work in all cases because there are explicit counterexamples. Nonetheless, the algorithms are widely used because the examples that make them fail do not seem to occur in practice. Exactly what constitutes "in practice" is again an empirical matter.

By the engineering component of numerical analysis I mean the design and implementation of numerical algorithms. These activities comprise at least half of the subject and perhaps more, as evidenced by the fact that in numerical analysis names tend to be associated with algorithms rather than theorems. Where topology gives us the Jordan curve theorem, Poincaré's conjecture, and Urysohn's lemma, numerical linear algebra gives us Jacobi's method, the Lanczos algorithm, and Householder transformations.

Algorithmic design is a skill quite different from discovering and proving theorems; it is more akin to inventing a gadget. What is required is a thorough knowledge of the parts from which algorithms may be assembled

and an intuitive appreciation of how they will behave on a computer. It is not surprising, then, that it is not necessary to be educated as a mathematician to build a good algorithm. For example, algorithms are at the heart of computer science, most of whose departments require deplorably little mathematical background from their students. Again, some of the best numerical algorithms come from scientists and engineers; for example, the finite element method was used by engineers long before numerical analysts took it up.

In establishing the first of the propositions stated at the beginning of this review we have gone far toward establishing the second: that a traditional mathematical education will not produce a numerical analyst. The empirical and engineering aspects are, quite properly, outside the province of the usual mathematics curriculum. Moreover, some of the most powerful theorems in pure mathematics become surprisingly weak when they are applied to numerical problems—the equivalence of norms in R^n is just one example.

One very important problem for the mathematically oriented is to appreciate the role of mathematical rigor in numerical analysis. As in all sciences, results in numerical analysis are frequently obtained by nonrigorous, intuitive modes of reasoning, with the justification that they can be tested empirically. Moreover, a premature attempt to be rigorous can be stultifying. Now in fact, this situation obtains in mathematics as well; a lot of nonrigorous reasoning goes into mathematical creation, reasoning which is then cleaned up in the final presentation. But many mathematicians have trouble transferring this creative process outside their specialties. They become mistrustful, and in their attempts to get things right they may earn reputations as pettifoggers or even obstructionists. The cure for this is to adopt the modes of reasoning appropriate to the discipline, while at the same time distinguishing what has and has not been rigorously established—a difficult, but not impossible balancing act. Unfortunately, the rather lax requirements of most universities make it easy for mathematics students to avoid any deep contact with other disciplines where this act could be learned.

Turning now to the third proposition, let us first note that there are a number of reasons why a pure mathematician might want to take up numerical analysis. It is an interesting and useful field with very tangible intellectual rewards, as anyone who has solved a difficult numerical problem knows. It is interdisciplinary, combining mathematics and computer science, usually with some area of applications. It abounds with unsolved research problems. On a more mundane level, anyone bailing out of a graduate seminar that has not attracted enough students will find that numerical analysis beats precalculus as a landing place. Finally, numerical analysts are in short supply and are eminently employable.

However, it is not easy for a mathematician to educate himself in numerical analysis. Perhaps the best way is to get a job solving numerical problems, and in the long run something like this is required to leaven theoretical knowledge acquired elsewhere. But for a variety of obvious reasons this approach is not open to many academic mathematicians. An alternative is to associate with numerical analysts, attending their seminars and discussing

problems. This option, however, requires that a generous number of numerical analysts be available locally.

Thus many mathematicians desiring an introduction to numerical analysis will have to get it by studying texts and monographs, supplemented by research papers. Regarding texts, the tyro is encouraged to seek expert advice; of the multitude of introductory books, a significant proportion are bad and all are deficient in one way or another. On the other hand there are a number of excellent monographs on a variety of topics in numerical analysis. Again, the help of an expert is advisable to avoid the outmoded and the eccentric.

The current definitive monograph on numerical linear algebra is J. H. Wilkinson's *The Algebraic Eigenvalue Problem*. Although it is by no means out of date, this fifteen-year-old book is beginning to show its age. It is therefore fortunate that Beresford Parlett has written this book on the symmetric eigenvalue problem. His aim is to produce a sequel to Chapter 5 of Wilkinson's book. He succeeds admirably.

There are at least three reasons why the symmetric eigenvalue problem is a good starting place for a mathematician to learn about numerical analysis. In the first place the underlying mathematical theory is nontrivial and elegant, so that the mathematician has a secure home base. Moreover, an unusually large proportion of the theorems are directly useful in computations. Second, the basic algorithms admit of clean analyses which are also sharp. Finally, the effects of rounding errors are more easily discerned than in the general case.

The book divides rather naturally into two parts. The first, consisting of chapters one through nine, is primarily concerned with small dense matrices. Much of this material is contained in Wilkinson's book; however, the pace is slower and more systematic. The first chapter treats the elementary theory. The very important second chapter introduces the practicalities that dominate matrix computations, especially the analysis of rounding error. The remaining six chapters are devoted to the basic algorithmic tools, concluding with a nice treatment of the ubiquitous QR algorithm. Chapter 9 scotches the notion that Jacobi's method is intrinsically simpler than the QR algorithm.

The last half of the book is directed toward large sparse matrices, and contains much recent, even new material. Chapter 10 treats perturbation theory for eigenvalues, and Chapter 11 the same for invariant subspaces, with particular attention being paid to Rayleigh-Ritz approximations. Chapter 12 lays the theoretical groundwork for the Lanczos algorithm, which is discussed in Chapter 13. These two chapters constitute the high point of the book, after which the chapters on subspace iteration and the generalized eigenvalue problem are somewhat anticlimactic. The latter chapter, incidently, is not up to the rest of the book, perhaps because theory and computational practice for definite generalized eigenvalue problems are less well developed.

The book is carefully written, and the author's intent is clearly to communicate all aspects of the subject. The exercises are well chosen and each chapter concludes with a useful notes and references section. There is a nice annotated bibliography at the end. One of the most important contributions of the book is to bring to the printed page some results of W. Kahan, which

would have otherwise languished unpublished in his notebooks.

The book is not without defects. The style, which is generally lively, sometimes degenerates into preciosity. The initial discussion of rounding error could be more leisurely. The author fails to warn the reader not to use the two-is-enough orthogonalization to compute a QR factorization unless he will be happy with “orthogonal” matrices having zero columns. What everyone else calls the generalized eigenvalue problem, the author calls the general linear eigenvalue problem. There are many minor technical points with which it would be easy to quibble.

But none of this alters the fact that Parlett has written a fine book. Because eigenvalue problems arise naturally in the analysis of vibrating systems, the author has collected some apt quotations about vibrations at the beginning of his preface. However, he missed the Beach Boys’ “I’m pickin’ up a good vibration”, which sums up my feelings about this book.

G. W. STEWART