Eigenvalue Computation in the 20th Century

Gene H. Golub* and Henk A. van der Vorst†

Abstract

This paper sketches the main research developments in the area of computational methods for eigenvalue problems during the 20th century. The earliest of such methods dates back to work of Jacobi in the middle of the nineteenth century. Since computing eigenvalues and vectors is essentially more complicated than solving linear systems, it is not surprising that highly significant developments in this area started with the introduction of electronic computers around 1950. In the early decades of this century, however, important theoretical developments had been made from which computational techniques could grow. Research in this area of numerical linear algebra is very active, since there is a heavy demand for solving complicated problems associated with stability and perturbation analysis for practical applications. For standard problems, powerful tools are available, but there still remain many open problems. It is the intention of this contribution to sketch the main developments of this century, especially as they relate to one another, and to give an impression of the state of the art at the turn of our century.

1 Sources

Numerical linear algebra is a very active field of research. Many problems are challenging of themselves, and in addition, much of scientific computing depends critically in one way or another on numerical linear algebra algorithms. Not only do the more classical scientific computational models for physical or engineering problems depend on linear algebra kernels, but many modern applications, such as information retrieval and image restoration, profit from numerical linear algebra results. These factors have motivated numerical linear algebra research throughout the entire twentieth century.

The field has blossomed, especially since the introduction of the modern computer, roughly from the early 1950's. This is evident from the large number of scientific journals in which articles in this area appear: *SIAM J. on Matrix Analysis and Applications (SIMAX)*, *Linear Algebra and its Applications (LAA)*, *Numerical Linear Algebra with Applications (NLAA)*, are completely devoted to this specialty. Articles on numerical linear algebra, theoretical as well as applied, regularly appear in journals such as *BIT*, *SIAM J. Numerical Analysis*, *SIAM J. on Scientific Computing*, *J. on Computational and Applied Mathematics*, *J. Applied Numerical Mathematics*, *Numerische Mathematik*, *Numerical Algorithms*, *Mathematics of Computation*, *Parallel Computing*, *ACM Transactions on Mathematical Software*, *Computing*, *J. Inst. Math. Applic.*, *SIAM Review*, *IMA J. Num. Anal.*, and several others in more application oriented directions, such as *J. Computational Physics* and engineering journals. And from, for instance, the bibliography in Golub and Van Loan's book, [51] one can see how many papers are referenced from these and other sources. A quick glance through the contents of the average 60 papers per year in SIMAX shows that roughly 40% of the papers are associated with eigenvalue problem research, and it is likely that this holds more or less for the many papers per year that focus on numerical linear algebra.

*SCCM, Stanford University, Stanford, USA, e-mail golub@sccm.stanford.edu. The work of this author was completed under a grant from the DOE: DE-FG03-97ER35330.

†Department of Mathematics, Utrecht University, Utrecht, The Netherlands, e-mail: vorst@math.uu.nl

This makes any attempt to write a complete overview on the research on computational aspects of the eigenvalue problem a hopeless task. It also serves as an excuse for the incompleteness in the current overview. We have tried to highlight what seemed most important from our point of view. We have included references to main sources, and we have made a personally colored selection of references to more specialized details. Instead of trying to give an overview of all sorts of different approaches that have been followed to solve aspects of eigenproblems, we will try to emphasize the history of those methods that, in our opinion, still play a role. Our aim is to consider the algorithmic developments from a historical point of view and to indicate how the recent powerful techniques are the result of many smaller steps. This will also help to show how many of the algorithms are interrelated; we hope not to get lost in sidesteps. The reader who is interested in methods that have played a role but that are at present no longer considered to be on the main track, is referred to Householder's and Wilkinson's books [64, 154]. In addition, Parlett [100] gives interesting historical information on older methods that still have some significance from a theoretical point of view.

In order to be active in this area of research, or to be informed about special aspects, then one might be interested in our main sources:

- Wilkinson: The Algebraic Eigenvalue Problem [154]
- Householder: The Theory of Matrices in Numerical Analysis [64]
- Wilkinson and Reinsch: The Handbook [158]
- Parlett: The Symmetric Eigenvalue Problem [100]
- Stewart and Sun: Matrix Perturbation Theory [129]
- Watkins: Fundamentals of Matrix Computations [150]
- Golub and Van Loan: Matrix Computations [51]
- Chatelin: Spectral Approximation of Linear Operators [18]
- Saad: Numerical Methods for Large Eigenvalue Problems [116]
- Demmel: Applied Numerical Linear Algebra [28]
- Trefethen and Bau: Numerical Linear Algebra [137]
- Arpack Guide (Lehoucq, Sorensen and Yang) [81]
- Dongarra *et al*: Numerical Linear Algebra for High Performance Computers [31]
- Bai's paper in Numerical Lin. Alg. with Appl. [6]
- Watkin's paper in SIAM Review [151]
- Dongarra and Walker on Software [33]
- Wilkinson: State of the Art overview [157]
- van der Vorst and Golub: State of the Art paper [50]

By regularly examining the dedicated numerical linear algebra journals, one should be able to trace most of the relevant and interesting papers for further investigations and research.

It should be noted that we have concentrated on algebraic eigenvalue problems in this paper. For eigenvalue problems related to, for instance, PDEs, one may use methods that exploit the nature of the PDE or the expected behavior of the solution. We have not considered such specialized techniques (of which multigrid is a good example).

2 Introduction

The eigenvalue problem for square matrices A, that is the determination of nontrivial solutions of $Ax = \lambda x$, is a central topic in numerical linear algebra. It is inherently nonlinear and this leads to many computational problems. Computation of the eigenvalues λ via the explicit construction of the characteristic equation

$$det(A - \lambda I) = 0$$

is, except for very special cases, not an option since the coefficients of the characteristic equation cannot be computed from determinant evaluations in a numerically stable way. And even if the characteristic equation could be determined accurately, then the computation of its roots, in finite precision, may be highly unstable since small perturbations in the coefficients may lead to large perturbations of the roots. The numerical computation of the associated eigenvectors and generalized eigenvectors is even more delicate, in particular when eigenvectors of A make small angles with each other. In the limiting case, when the matrix is defective, A can be reduced to the Jordan canonical form, but arbitrary small perturbations in A may yield a nondefective matrix. This leads to many challenging numerical questions, which give rise to the central problem: how can we compute eigenvalues and eigenvectors in an efficient manner and how accurate are they?

In fact, this was already recognized by Jacobi, who, in 1846, computed the eigenvalues of symmetric matrices by rotating the matrix to a strongly diagonally dominant one. We will return to this later, since Jacobi's techniques are still relevant and have led to popular and powerful algorithms.

Another longstanding method that is of great significance and serves as the basis for many algorithms is the Power iteration. The method is based on the idea that if a given vector is repeatedly applied to a matrix, and is properly normalized, then ultimately, it will lie in the direction of the eigenvector associated with the eigenvalues which are largest in absolute value. The rate of convergence for the Power iteration depends on the ratio of the second largest eigenvalue (in absolute value) to the largest eigenvalue (in absolute value) and for many applications this leads to unacceptably slow convergence. The method can be problematic if one wants to compute a number of extremal eigenvalues. The Power iteration is still in use, but most frequently as (implicit) part of more efficient techniques, e.g. Krylov methods, inverse iteration, QR-method.

What becomes clear is that all these methods are of an iterative nature, and this is necessarily the case, since if there were a method of computing the eigenvalues of an nth order matrix in a finite number of computations, depending only on n, then this would be in contradiction with the fundamental theorem of Abel-Ruffini (and also a well-known result in Galois theory) that no such algorithm exists for the computation of the roots of a general polynomial of degree greater than 4. Hence, an algorithm for a matrix with a general structure (that is, not a diagonal matrix or a triangular matrix or alike) is necessarily iterative and the problem is to identify iterative algorithms which have a fast rate of convergence and lead to accurate results.

In solving an eigenvalue problem there are a number of properties that need be considered. These greatly affect the choice of algorithm. We list below a number of questions that an investigator needs to consider in solving a particular problem.

- Is the matrix real or complex?
- What special properties does the matrix have?
  - symmetric
  - Hermitian
  - skew symmetric
  - unitary
- Structure?

  - band
  - sparse
  - structured sparseness
  - Toeplitz
- Eigenvalues required?
  - largest
  - smallest in magnitude
  - real part of eigenvalues negative
  - sums of intermediate eigenvalues

As well as the standard eigenproblem, there are a variety of more complicated eigenproblems, for instance $Ax = \lambda Bx$, and more generalized eigenproblems like $Ax + \lambda Bx + \lambda^2 Cx = 0$, higher order polynomial problems, and nonlinear eigenproblems. All these problems are considerably more complicated than the standard eigenproblem, depending on the operators involved. However, as the standard eigenproblem has become better understood, in a numerical sense, progress has been made in the other problems and we will consider developments in solving these problems.

3 Canonical Forms

The standard approach for the numerical solution of the eigenproblem is to reduce the operators involved to some simpler form, that yields the eigenvalues and eigenvectors directly, for instance, diagonal form. The idea is that the transformation be made with orthogonal operators as often as possible, in order to reduce the effects of perturbations.

The easiest situation is the symmetric case: for a real symmetric matrix, there exists an orthogonal matrix Q, so that $Q^T AQ = D$, where D is a diagonal matrix. The diagonal elements of D are the eigenvalues of A, the columns of Q are the corresponding eigenvectors of A.

Unsymmetric matrices do not in general have an orthonormal set of eigenvectors, and may not have a complete set of eigenvectors, but they can be transformed unitarily to Schur form:

$$Q^* AQ = R,$$

in which R is upper triangular. In fact, the symmetric case is a special case of this Schur decomposition, since a symmetric triangular matrix is clearly diagonal. Apart from the ordering of the eigenvalues along the diagonal of R and the sign of each column of Q, the matrix Q is unique. Van Dooren [146] has described an algorithm for the orthogonal transformation of Q, so that the eigenvalues appear in prescribed order along the diagonal of R. If the eigenvalues are distinct then there exists a nonsingular matrix X (in general not orthogonal), that transforms A to diagonal form:

$$X^{-1} AX = D.$$

An unsymmetric matrix can be transformed to Jordan form by a nonsingular matrix X. This Jordan matrix may have upper bidiagonal blocks along the diagonal. Each of these blocks has identical eigenvalues and the upper bidiagonal elements are equal, and most often set to 1. The numerical computation of the Jordan form is highly unstable, since a small perturbation suffices to obtain a matrix with different eigenvalues (and possibly a complete eigenvector system). Small angles between (some of) the eigenvectors reveal that A is close to a matrix that is similar to a non-diagonal Jordan form. For a discussion on how to compute elementary Jordan blocks (with the help of the Singular Value Decomposition), see the 1976 paper by Golub and Wilkinson [52].

Just as the Jordan canonical form describes the eigenstructure of a matrix, the Kronecker form does this for matrix pencil $A - \lambda B$, even for rectangular A and B. For details on this we refer to papers by Wilkinson [155, 156, 143], and Kågström [66]. The latter has also developed software for the computation of the Kronecker structure [67]. Wilkinson, in his discussion on the progress made in the period 1976-1986 in eigenvalue computations [157], noted that the Jordan canonical and Kronecker canonical forms were largely regarded as irrelevant by numerical analysts because of their ill-posedness. He even stated: "*Many felt that I should have ignored the Jordan canonical form in the Algebraic Eigenvalue Problem [154] and I had misgivings about including a discussion of it*". Since the seventies, this has changed, and contributions have been made by many, including Demmel, Beelen, Van Dooren, Chaitin-Chatelin, Edelman, Kågström, Nichols, Kautsky, Golub, and Wilkinson. Although serious attempts have been undertaken for the computation of the Kronecker canonical form, by for instance Kågström and Van Dooren, this still needs further research. Also the computation of invariant subspaces of highly nonnormal matrices is still in its infancy, notwithstanding useful contributions by, for instance, Chaitin-Chatelin *et al* [17, 12] and Lee [80]. For recent references, see [4]. Van Dooren described, in papers published in 1981, how the Kronecker form can be used in system control problems (input-output systems) [145, 144].

Related to eigendecompositions is the singular value decomposition. Let A be a real m by n matrix, then there exists an orthogonal m by m matrix U and an orthogonal n by n matrix V, such that

$$U^T AV = diag(\sigma_1, \ldots, \sigma_p),$$

with $p = min\{m, n\}$, and $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_p \geq 0$. The values σ_i are the *singular values* of A, the columns v_i of V are the *right singular vectors* and the columns u_i of U are the *left singular vectors*. The number of nonzero singular values is equal to the rank of the matrix A. The singular value decomposition (SVD) plays an important role in numerical linear algebra, for instance in the solution of underdetermined or overdetermined linear systems.

4 Perturbation theorems

Perturbation theorems play a very essential role in computational processes for eigenproblems. As we have remarked, eigensolvers are essentially iterative processes and many of them rest on the principle of reducing a matrix to a special/simpler form, either diagonal or upper triangular. One has then to decide when a matrix is sufficiently close to the appropriate ultimate form. It is important to know what the approximation errors imply about the desired eigen information. A modern treatment of perturbation theory for a variety of eigenproblems is given in the book by Stewart and Sun [129]. We will restrict ourselves here to what we regard as some of the most relevant results in this area.

Many theoretical results rest on the famous Gershgorin Disc Theorem, which states that the eigenvalues of a matrix $A = (a_{ij})$ are located in the union of circles with centre a_{ii} and radius $\sum_{j \neq i} |a_{ij}|$. This theorem first appeared in a classic paper by Gerschgorin in 1931 [44]. A very useful refinement shows that if a group of s discs D_i is isolated from the others, then this group contains precisely s eigenvalues. In particular, if one disc is isolated then this disc contains one eigenvalue. This particular case is of great interest, since it can be used for stopping criteria in actual computations.

Wilkinson [154] discussed the application of Gerschgorin's theorem to various situations. He mentioned the discussion of useful extensions by Taussky in her 1949 paper [131]: Varga acknowledged Taussky's paper in his work on solving systems of linear equations by iterative methods (cf. [148]). An important extension, using block matrices, of the Gerschgorin Theorem was given by Feingold and Varga [38].

The eigenvalues depend continuously on the elements of A and if the ith eigenvalue λ_i is distinct, then it is even differentiable. In this case one can carry through a first order perturbation analysis (cf. [129, p.185]). This leads to the observation that if a matrix A is perturbed by δA, then the perturbation to λ_i is in first order of terms of δA given by

$$\delta\lambda_i = \frac{1}{y_i^* x_i} y_i^* \delta A x_i,$$

where x_i, and y_i are the normalized right and left eigenvectors respectively, corresponding to λ_i, and y_i^* denotes the complex conjugate of y_i. The factor $\frac{1}{y_i^* x_i}$ is referred to as the *condition number* of the ith eigenvalue. The Bauer-Fike result (1960) [9], which is actually one of the more famous refinements of Gershgorin's theorem, makes this more precise: the eigenvalues $\widetilde{\lambda}_j$ of $A + \delta A$ lie in discs B_i with center λ_i, and radius $n\frac{\|\delta A\|_2}{|y_i^* x_i|}$ (for normalized x_i and y_i).

The Courant-Fischer minimax theorem is the basis of many useful results. For a symmetric matrix A with ordered eigenvalues $\lambda_n \leq \ldots \leq \lambda_2 \leq \lambda_1$ it states that the eigenvalues are the stationary values of the *Rayleigh quotients*:

$$\lambda_k = \max_{dim(S)=k} \min_{0 \neq y \in S} \frac{y^* A y}{y^* y},$$

for $k = 1, 2, \ldots, n$. Some important consequences are the following.

For symmetric matrices, Weyl in 1912 proved an important property for symmetric perturbations δA:

$$|\widetilde{\lambda}_i - \lambda_i| \leq \|\delta A\|_2.$$

In fact, Weyl gave even more detailed results in terms of the eigenvalues of A and of δA: let the eigenvalues of δA be denoted by $\eta_n \leq \ldots \leq \eta_2 \leq \eta_1$, then the eigenvalues $\widetilde{\lambda}_i$ of $A + E$ satisfy

$$\lambda_k + \eta_n \leq \widetilde{\lambda}_k \leq \lambda_k + \eta_1.$$

These inclusion formulas were later refined to, in Parlett's terminology, a blizzard of results, independently obtained by Kahan (1957)[1] and Weinberger (1974) [153]. See Parlett's book [100, Ch.10.6] for an exposé of these results. It is interesting to note that by 1990 the theory had evolved to such an extent that Weyl's result and the Kato-Temple results could be left as exercises in Stewart and Sun's book [129, p.210-211]. This illustrates a rich and powerful framework.

Another important property that plays a big role in iterative (projection type) algorithms, is the interlace property. Let A_r denote the leading r by r minor of A, with eigenvalues $\lambda_j^{(r)}$, then

$$\lambda_{r+1}^{(r+1)} \leq \lambda_r^{(r)} \leq \lambda_r^{(r+1)} \leq \cdots \leq \lambda_2^{(r+1)} \leq \lambda_1^{(r)} \leq \lambda_1^{(r+1)}.$$

An important result, that underlies the powerful *divide and conquer* method, comes from rank-one perturbations. If $B = A + \tau cc^T$, with $\|c\|_2 = 1$, and real $\tau \geq 0$, then the ith eigenvalue of B is in the interval $\lambda_i, \lambda_{i-1}]$, for $\tau \leq 0$ it is in $[\lambda_{i+1}, \lambda_i]$. In either case, there exist nonnegative $\alpha_1, \ldots, \alpha_n$ with $\sum_i \alpha_i = 1$, such that $\lambda_i(B) = \lambda_i + \alpha_i \tau$.

Further details and results can be found in most books on the (numerical) eigenproblem; in particular Wilkinson's book [154] is a great source. A good overview of these results and similar results for invariant subspaces is given in [51]. From a result formulated as the Kato-Temple theorem, one can obtain sharp bounds for the Rayleigh quotients for symmetric matrices. This rests on work of Temple (1933) [132] and Kato (1949) [70]; more extended work in this direction has been done by Davis and Kahan (1970) [27] (see also [18, p.46], [28, Ch.5.2]). In these results the *gap* for the ith eigenvalue plays an essential role: $gap(i, A) \equiv \min_{j \neq i} |\lambda_j - \lambda_i|$. A small gap indicates a sensitive eigenvector. In particular, let x denote a normalized vector with Rayleigh

[1] according to Parlett [100, p.203], who also mentions unpublished (?) results of Weinberger (1959)

quotient $\beta = x^T Ax$, and residual $r = Ax - \beta x$. Then there is a (normalized) eigenvector q_i, corresponding to λ_i, for which

$$|\lambda_i - \beta| \leq \frac{\|r\|_2^2}{gap(i, A)}, \quad \sin(\theta) \leq \frac{\|r\|_2}{gap(i, A)},$$

where θ denotes the angle between x and q_i. These results show the superior quality of a Rayleigh quotient from a given subspace. It is exploited in modern iterative methods, such as the Lanczos algorithm, but it is also essential in the QR algorithm. Related to the perturbation analysis for Rayleigh quotients is work of Kaniel (1966) [69] for errors in the Ritz approximations computed in the Lanczos process. For a comprehensive discussion of this, see [100, Ch.12].

5 Jacobi's method

For our discussion of the Jacobi method, we have used the following sources: [154, Ch.5], [100, Ch.9], and [51, Ch.8.4]. The Jacobi method which was originally proposed in 1846 [65], reduces a real symmetric matrix to diagonal form by a sequence of plane rotations. Jacobi, however, did not use the method to full convergence, as was done in the twentieth century, but combined it with an iterative process (Gauss-Jacobi iterations) for the missing component of the desired eigenvector (for which he took a unit vector as an initial guess). Actually, Jacobi's technique may be viewed as a form of preconditioning for the Gauss-Jacobi iteration that he also used to solve linear least squares systems. This has escaped the attention of most researchers that were active with Jacobi's method; the exception seems to be Bodewig [2, p.280-287]. The preconditioning part of the method, as an iterative technique to diagonalize a matrix, was reinvented in 1949 by Goldstine, Murray, and von Neumann and published in a manuscript. After Ostrowski had pointed out that this was actually a rediscovery of Jacobi's method[2], the adapted manuscript was published only in 1959 [46][3]. According to Wilkinson [154, p.343] Jacobi's method was already being used independently on desk computers at the National Physical Laboratory in 1947. From 1950 on the method got much attention. In the classical process, the maximal off-diagonal element is annihilated, and this guarantees convergence. Since it is a time-consuming process to determine the maximal element after each rotation, cyclic procedures were suggested, (cf.[54]). Later, threshold strategies were developed in order to avoid wasting time in trying to annihilate tiny elements [105]. Quadratic convergence for the cyclic Jacobi algorithm was proven, under various assumptions, by Henrici (1958) [59], Schönhage (1961) [119], Wilkinson (1962) [154], and van Kempen (1966) [147]. This rate of convergence sets in after a number of sweeps (that is $(n-1)n/2$ elementary rotations), but there is no rigorous bound on the number of sweeps required to achieve a specified accuracy. Brent and Luk [14] argued that this number is proportional to $\log(n)$, which is in line with Parlett's remark [100, p.181] that after three or four sweeps through all the off-diagonal elements convergence is usually very rapid.

The success of the Jacobi method for diagonalizing a symmetric matrix by orthogonal similarity transformations inspired many investigators to find a similar method for non-symmetric matrices. It was quickly realized that the Schur form was the appropriate decomposition. John Greenstadt, in 1955 [53], was one of the earliest investigators to develop such a method (indeed, Greenstadt made von Neumann aware of this canonical form). Unfortunately, these earliest attempts were not successful. The QR method, that gained more popularity somewhat later, can be viewed, however, as a Jacobi like method, since it can produce the Schur form via a sequence of similarity transformations composed of rotations.

[2]Michele Benzi brought to our attention that this story is narrated on page 294 in Goldstine's book *The Computer from Pascal to von Neumann*, Princeton University Press, 1972; it is also nicely covered in the Master Thesis of Anjet den Boer [30] on the history of Jacobi's method

[3]Parlett [100, p.184] dates the reinvention in 1946, by Bargmann, Montgomery, and von Neumann [7], but this is presumably a misprint

Rutishauser made an Algol60 implementation of Jacobi's process, as a contribution to the Wilkinson Reinsch collection [158]. In the sixties, the popularity of the method declined, because of the growing popularity first of the Givens method and slightly later, the Householder method: these latter two methods first reduced the matrix to tridiagonal form and then used an efficient procedure for computing the eigenvalues of the tridiagonal matrix. Interest in the Jacobi returned with the advent of parallel computers, starting with a paper of Sameh (1971) [118], and followed by others in the eighties. Variants of Jacobi's method were proposed; we mention the extensions for normal matrices, by Paardekooper in 1971 [96], and to nonnormal matrices by Eberlein, in 1970 [36]. The latter process was also part of the Wilkinson and Reinsch collection (the Algol60 procedure **eigen**). In order to improve data locality for distributed memory computers, block Jacobi methods were suggested, see Bischof [10] for a discussion on the influence of the solution of the subproblems on the overall process.

Another interesting feature of the Jacobi method is its superior behaviour with respect to accuracy. Wilkinson analysed this and showed that the relative error in the eigenvalue approximations is eventually reduced to the order of the condition number of A times machine precision. This was perfected in 1992, by Demmel and Veselič [29], who showed that for symmetric positive definite matrices, the condition number of A could be replaced by that of the matrix symmetrically scaled by the diagonal. If one is satisfied with less accuracy, then for large scale computations, Jacobi's method is no longer regarded as competitive, not even for modern parallel computers.

6 Power method

For our discussion of the Power Method, we have borrowed material from Householder's book [64]. The Power Method, for general square matrices, is the simplest of all the methods for solving for eigenvalues and eigenvectors. The basic idea is to multiply the matrix A repeatedly by a well chosen starting vector, so that the component of that vector in the direction of the eigenvector with largest eigenvalue in absolute value is magnified relative to the other components. Householder called this *Simple Iteration*, and attributed the first treatment of it to Müntz (1913). Bodewig [2, p.250] attributes the power method to von Mises [149], and acknowledges Müntz for computing approximate eigenvalues from quotients of minors of the *explicitly* computed matrix A^k, for increasing values of k. For a careful analytic treatment of the Power Method, Householder acknowledged work by Ostrowski and Werner Gautschi; the reader can find a fairly complete treatment in Wilkinson's book [154] together with the proper references. The speed of convergence of the Power iteration depends on the ratio of the second largest eigenvalue (in absolute value) to the largest eigenvalue (in absolute value). In many applications this ratio can be close to 1 – this has motivated research to improve the efficiency of the Power Method. It is interesting that the most effective variant is the inverse Power Method, in which one works with the matrix $(A - \mu I)^{-1}$, and this variant was proposed as late as 1944 by Wielandt (Wielandt's *fractional iteration*). Wielandt also proposed continuing the process after the largest eigenvalue has converged, by working with the deflated matrix $A - \lambda vv^*$, for which λ, v is the computed eigenpair (with $\|v\|_2 = 1$), associated with the largest eigenvalue in magnitude. (The deflation procedure outlined here is for symmetric matrices. For unsymmetric matrices it is necessary to work with at least two vectors; the choice of one of the vectors may not be unique.) This is called *implicit deflation*; another possibility is to keep the iteration vectors orthogonal to the computed eigenvector(s): *explicit deflation*. A compact description and analysis of these deflation techniques was given by Parlett [100]. The Power Method and the Inverse Power Method, in their pure form are no longer competitive methods even for the computation of a few eigenpairs, but they are still of interest since they are explicitly or implicitly part of most modern methods such as the QR method, and the methods of Lanczos and Arnoldi. These methods evolved in some way or another from the Power Method and some of the techniques that were suggested as improvements to the Power Method are still in use as acceleration techniques for modern iterative methods. One of these ideas is to work with polynomials of A, with the purpose of damping unwanted parts of the spectrum.

Another possibility is working with properly updated shifts μ in the inverse process and, in particular, if one takes the Rayleigh quotient with the most recent vector as a shift, then one obtains the *Rayleigh quotient iteration*. According to Parlett [100, p.71], Lord Rayleigh used in the 1870s a less powerful technique: he did a single shift-and-invert step with a Rayleigh quotient for an eigenvector approximation, but with a unit vector as the right-hand side. (This saves the refactoring of the matrix $(A - \mu I)$ at each iteration). The modern RQI, in which one takes the most current eigenvector approximation as the right hand side, leads to very fast convergence. Ostrowski, in a series of 6 papers [95], studied the convergence properties for variance of RQI for the symmetric and unsymmetric case. He was able to establish cubic convergence in both cases under various circumstances (in the unsymmetric case for a properly generalized Rayleigh quotient). These results are essential for the understanding of modern iterative techniques that are based on (approximate) shift-and-invert strategies (for example, the Jacobi-Davidson method, see below).

A step forward was to work with a set of independent vectors in order to find a number of eigenvectors, instead of the deflation procedure suggested by Wielandt. A problem with the Power Method is the determination of eigenvalues that have equal modulus, for instance, finding a conjugate pair of eigenvalues of a real unsymmetric matrix. It is therefore quite natural to work with a couple of independent vectors: this was first suggested in 1937 by Horst [62]. The next step that seems logical, in hindsight, is to force the vectors to be independent. This was initially done (cheaply) by Gaussian tranformations by Bauer [8] in 1957, and led to *Treppeniteration*. If the set of vectors is denoted as a matrix L_s (an n by s unit lower-triangular matrix), then one forms AL_s and factors the resulting matrix, by Gaussian transformations, as $L_{s+1}R_{s+1}$. If the eigenvalues are distinct, then the s by s upper triangular matrix R_{s+1} converges, for increasing s, to a matrix whose eigenvalues are those of the dominant subspace on its diagonal. Rutishauser [111] made the important observation that if we factor A as $A = LR$ (again L unit lower triangular), then the similar matrix $L^{-1}AL = L^{-1}LRL = RL$. He proposed decomposing RL again, and repeating this process in an iterative fashion. This R also converges to an upper triangular matrix, and L is a unit matrix. This is the *LR method* of Rutishauser. The correspondence between Treppeniteration and LR is that if we start Treppeniteration with a unit full matrix, then in exact arithmetic we obtain the same matrices R in the process. For an efficient implementation, the matrix A is first reduced to an upper Hessenberg matrix. The LR method maintains this form throughout the process, and this makes LR computationally very attractive. Rutishauser's observation that permuting the factors of the matrix is equivalent to performing a similarity transformation was a key step. Wilkinson [154, p.485] commented: *"In my opinion its development (i.e. of LR) is the most significant advance which has been made in connection with the eigenvalue problem since the advent of automatic computers"*. However, Bauer's technique could be applied to a smaller set of starting vectors and it does not modify the matrix A. For this reason, in the words of Householder [64, p.195], it is self-correcting. This seems to imply that Treppeniteration leads to more accurate results.

Since orthogonal reduction techniques often evidence superior stability properties, it became apparent that the LR factorization should be replaced by a QR factorization. This leads to one of the most popular and powerful methods of our time for eigenvalue problems: the QR method for computing all of the eigenvalues and associated eigenvectors of a dense symmetric matrix. (In fact, the QR method has essential enhancements that make the method really powerful; we will discuss this in another section).

With the number of vectors less than n, this Power Method in combination with QR orthogonalization is known as the *Simultaneous Iteration method*; Rutishauser studied this method in 1969 [112], see also [113]. Its convergence behaviour for general unsymmetric matrices was studied by Stewart [126] in 1976. Stewart also developed a subroutine, based on simultaneous iteration, for the computation of a dominant invariant subspace. This routine, SRRIT [127], was further improved in 1992, and made available for general use through Netlib [5].

The collection of vectors generated by the Power Method define Krylov subspaces of increasing dimension. This motivated Krylov to try to determine the characteristic polynomial of a matrix by inspecting the dependence of a full set of these vectors. This procedure may fail because the system of equations is highly ill-conditioned but this can be repaired by orthogonalizing each new vector to the previous vectors and applying A onto the last constructed vector. This iteration process is known as the Lanczos method for symmetric matrices, and Arnoldi's method for unsymmetric matrices. We will discuss these Krylov methods below.

Our presentation might suggest that the Krylov methods have overshadowed the Simultaneous Iteration method, and for most situations this is indeed the case. Parlett, however, [100, p.289] described situations where Simultaneous Iteration is still competitive. For instance, if we can store only a limited number of n-vectors in fast memory, or if the relative gap between the desired eigenvalues and the others is great, then Simultaneous Iteration is very useful.

7 Reduction algorithms

Early computational techniques, other than Jacobi's famous but slowly converging diagonalization method, and the unsatisfactory Power method with its many restrictions, attempted to exploit the fact that every matrix satisfies its characteristic equation. To this end, Krylov suggested in 1931 [73], using the vectors x, Ax, A^2x, ..., generated by the Power method, to determine the coefficients of the characteristic equation. This was not successful, because, as we have learned from Wilkinson's analysis [154] the roots of a polynomial may vary widely with only tiny perturbations to the coefficients of the polynomial. Even rounding the exact coefficients in floating point arithmetic may destroy much accuracy in many of the roots. Although Krylov's method failed, his name is still attached to the subspace generated by the Power method.

There is yet another reason for the failure of Krylov's method in finite precision arithmetic: the vectors generated by the Power method tend to converge in the direction of the eigenvectors associated with the dominating eigenvalues. Hence, the computed vectors for the subspace necessarily yield a very ill-conditioned basis. Checking mutual dependence of this basis, as is required in order to construct the characteristic polynomial, is an almost impossible task.

An early attempt to reduce the matrix A to a form that lends itself better for solving the characteristic equation was suggested by Hessenberg [60]. He suggested to compute a modified Krylov basis by making a set of basis vectors for the Krylov subspace, orthogonal to a given test-set, for instance the canonical basis vectors. This led to a reduction of A to upper Hessenberg form. This technique is very close to the techniques by Lanczos and Arnoldi.

In 1950, Lanczos [78] suggested building a basis for the Krylov subspace in a more stable way by orthogonalizing the vectors as they are constructed. The idea was to immediately orthogonalize the new vector for the Krylov subspace with respect to the already existing orthonormal basis. The difference with Hessenberg's approach is that Lanczos (and slightly later also Arnoldi) took the Krylov vectors themselves for the test-set.

The new vector for expansion is created by applying A to the latest orthogonal basis vector. For symmetric matrices A, this leads to a three-term recurrence relation between the basis vectors v_j and in exact arithmetic this can be formulated as

$$AV_j = V_jT_j + \gamma_j v_{j+1} e_{j+1}^T.$$

Obviously, this recursion must terminate for some $j \leq n$, in which case V_j forms the basis for an invariant subspace of A, and the eigenvalues of the tridiagonal matrix T_j are the eigenvalues of A with respect to this invariant subspace. This algorithm is equivalent to the well-known algorithm of the Dutch mathematician Stieltjes for generating orthogonal polynomials by a three-term recurrence relationship. Lanczos also proposed a reduction process for unsymmetric matrices A, the so-called two-sided Lanczos process. In this process two sets of basis vectors are constructed, one for the Krylov subspace with A and one for a Krylov subspace with A^T. By requiring bi-orthogonality of the two sets, the two bases can be used for reduction of A to tridiagonal form. This form has suffered from many drawbacks. Not only are the reduction matrices non-orthogonal, a suspect property, but the algorithm also suffers from various break-down conditions. (The basic problem lies in the fact that the measure generated by the initial vectors is not non-negative). The symmetric variant did not become popular in the 1950s, since it was soon recognized that rounding errors could spoil the process dramatically. Wilkinson [154] showed that the Lanczos algorithm is highly (forward) unstable and there seemed to be no way of stabilizing the process other than re-orthogonalizing the generated vectors. He showed this process is comparable to the methods of Householder or Givens (proposed in the late fifties), but the latter are more economical. He then concluded: "*it is difficult to think of any reason why we should use Lanczos' method in preference to Householder's.*" This illustrates that the Lanczos method was commonly viewed as a direct reduction method at that time, and from that point of view Wilkinson's remarks were quite correct.

At about the same time as Lanczos, Arnoldi (1951) [1] gave a reduction algorithm for unsymmetric matrices. This was basically the same algorithm as Lanczos' algorithm for symmetric matrices, with the difference that each new basis vector had to be orthogonalized with respect to all previous basis vectors. In this way A is reduced by an orthogonal similarity transformation to upper Hessenberg form. Arnoldi's method suffers far less from numerical instability, depending on how well the orthogonalization process is carried out. But the method is more expensive than the Householder reduction, making it less attractive, as a direct method, for large (dense) matrices.

A very important notion was the recognition that matrices could be reduced, by orthogonal transformations, in a finite number of steps, to some special reduced form that lends itself more efficiently to further computations. In particular, a symmetric matrix can be reduced to tridiagonal form by Jacobi-rotations, provided that these rotations are restricted to annihilate entries of A outside its tridiagonal part. This was suggested by Givens in 1954 [45], and in this connection the Jacobi-rotations are also called Givens rotations. A few years later, Householder, in 1958 [63], discovered that complete columns of A could be reduced to zero, outside the tridiagonal part, by the more efficient Householder reflections. These are well-chosen orthogonal rank-one updates of the form

$$(I - \frac{2}{v^T v} vv^T);$$

these are discussed in [139]. The Householder method has become the method of choice for the reduction of matrices to tridiagonal form on serial computers. Thus for eigenproblems, a symmetric matrix can be reduced by orthogonal similarity transformations to tridiagonal form and unsymmetric matrices can be transformed to upper Hessenberg form.

By 1960, the eigenvalue problem for a symmetric tridiagonal matrix was solved by using the Sturm sequence property for successive subdeterminants[4]. The corresponding eigenvectors were solved by inverse iteration. The whole process is described in Givens' papers. A complete and thorough analysis for the Givens and Householder reductions and for the use of the Sturm sequences, is given in Wilkinson's book, which was the numerical linear algebra bible (Old Testament) for a long time.

As we have already shown in the section on the Power Method, the QR method is, for determining the complete set of eigenvalues and eigenvectors, a superior technique. At the time that Wilkinson's book appeared, the blossoming of the QR method had just begun. Wilkinson devoted much attention to this method, but not as the method of choice for symmetric problems. We quote from Parlett [100, p.172]: "*Yet it was not invented until 1958-1959 and was not appreciated until the mid-1960s. The key idea came from Rutishauser with his construction of a related algorithm called LR in 1958.*" Whereas Wilkinson's book was the reference for eigenvalue problems in the period 1960-1980, after 1980, Parlett's book *The Symmetric Eigenvalue Problem* became the main

[4] This shows that working with the characteristic equation, if not explicitly constructed, is not a bad idea in some situations.

source, at least for symmetric problems. Comparison of the two books clearly shows the progress made in this field.

The use of the QR (a mathematical equivalent is the QL algorithm) algorithm began with the work of Francis, [40] who recognized in 1961-1962 that a QR iteration maps a Hessenberg matrix to a Hessenberg matrix again, and this makes the process economical and also adds to stability since the zero elements need not be computed. Furthermore, Francis cleverly implicitly used origin shifts, and these can be carried out very economically for Hessenberg matrices. Kublanovskaja, in 1961 [74], also independently discovered the same process, but did not employ the invariance of the Hessenberg form. She deeply understood the mathematical aspects of the algorithm but was less concerned with the important computational details. The inclusion of Wilkinson shifts eventually makes the process very efficient, and for these shifts it can be proved that, for symmetric matrices, the process does not fail. The order of convergence for symmetric matrices is cubic (see, for instance, [61]), while for unsymmetric matrices it is quadratic [51, 28]. These results rest on work of Ostrowski carried out in connection with the shift-and-inverse Power Method (the RQI method, see that section). For a treatment of modern implementations of the QR method see [51, Ch.7] or Demmel's book [28, sect.4.4.5]. These implementations incorporate techniques developed in the 1990s, such as (multiple) implicit shifts. This implicit shift technique leads to a rank-one perturbation of the Hessenberg structure, and this perturbation can be removed in its turn by a technique that is known as *chasing the bulge*: the perturbation (bulge) is chased down (and out of) the Hessenberg matrix with (double) Givens transformations. These chasing techniques were analysed in 1991 by Watkins and Elsner [152]. An important and complete overview of the practical QR algorithm can be found in [151].

With respect to the relation between the QR and the Lanczos algorithms, we note the following. Lanczos' method focusses on one particular starting vector; QR starts with a full orthogonal basis and keeps it orthogonal through the Power iterations; the inclusion of shifts does not destroy the structure of the Hessenberg matrix. With Lanczos' method, a shift only makes sense in damping unwanted parts of the spectrum, but one cannot vary the shift during the process.

By 1970 the standard numerical eigenproblem, for dense matrices of not too large order, could be regarded as essentially solved and research shifted to larger problems and other eigenproblems. The next important problem to consider was the generalized eigenproblem $Ax - \lambda Bx = 0$. An obvious approach is to transform this to a standard eigenproblem by inverting either A or B, or to work with more complicated transformations, such as the *Cayley Transform*: $(A-\sigma B)^{-1}(A-\tau B)$. These approaches share the disadvantage that matrices A and B are not treated in the same way, which is most obvious from the simple transformation $B^{-1}A$. This leads to problems if B is singular or ill-conditioned, but as Stewart (1980) has pointed out, this does not necessarily mean that the given eigenproblem is ill-conditioned because if A is well-conditioned then the pencil $B - \mu A$ may be well behaved (small perturbations in A and B lead to small perturbations in μ. In Stewart's analysis the matrices are treated symmetrically; in particular he suggests considering the pencil $\sigma A - \tau B$, and regarding multiples of (σ, τ), for which the determinant $|\sigma A - \tau B|$ vanishes, as generalized eigenpairs. See [129, Ch.VI] for more details. This approach, of course, requires a different reduction and this is accomplished by the Generalized Schur Decomposition, proposed by Moler and Stewart in 1973. This says that for arbitrary square A and B there exist unitary Q and Z such that $Q^*AZ = T$ and $Q^*BZ = S$ are upper triangular. For real matrices, the arithmetic can be kept real, but then the reduced matrices are quasi-triangular (that is 2 by 2 non-zero blocks along the diagonal may occur). Moler and Stewart [91] also proposed a stable algorithm to accomplish the reduction to (quasi) triangular form and this is known as the *QZ algorithm*. Major modern software packages include software for the QZ algorithm. For perturbation analysis, we refer to Stewart and Sun [129].

After the 1970s, the eigenproblem for dense matrices of moderate order seemed to be solved and further improvements were not expected, especially for symmetric dense matrices. However, with the ever increasing demand for higher efficiency and/or better accuracy, things changed from time to time. In 1981 Cuppen [25] proposed a divide and conquer algorithm for the solution of the eigenproblem for symmetric tridiagonal matrices. The idea was to split the tridiagonal matrix in two blocks, each of half the original size, plus a rank-one update. Cuppen showed how the eigenproblems for each of the blocks could be combined for the original full problem by exploiting the rank-one update property, which led to the solution of a secular equation. Initially, this approach was not seen as a competitive algorithm by itself for general matrices of modest dimensions, although Cuppen recognized that his algorithm was asymptotically (much) faster than QR. Further investigations by others were made on account of promising parallel properties. A major problem was that the original algorithm suffered from instabilities, especially for the eigenvectors belonging to close eigenvalues. Some scaling problems were rectified by Dongarra and Sorensen in 1987 [32], but the "right" implementation, according to Demmel [28, Sec.5.3.3] was not discovered until 1992 and published in 1995, by Gu and Eisenstat [55]. Meanwhile, software for this algorithm found its way into LAPACK and ScaLAPACK. As stated by Demmel again, the divide and conquer approach is now the fastest algorithm for computing all eigenvalues and eigenvectors of a symmetric matrix of order larger than 25; this also holds true for non-parallel computers. If the subblocks are of order greater thant 25, then they are further reduced; else, the QR algorithm is used for computing the eigenvalues and eigenvectors of the subblock. For a full treatment of the modern variant of the divide and conquer method, we refer to Demmel's book [28]. A recent discussion on parallel implementation aspects of this method can be found in [134].

There are still niches for other methods for dense symmetric matrices. Wilkinson advocated the bisection method for tridiagonal matrices if only a small subset of the eigenvalues is wanted. Inverse iteration may then be used to determine the corresponding eigenvectors. For the bisection method, based on the Sturm sequence property, the reader may find classical material in [154], a modern treatment of the inverse iteration (considered suspect by many because of the ill-conditioning of the shifted matrix) can be found in [100].

8 Iterative Methods

Soon after its introduction in 1952, it was recognized that the Lanczos method was not a panacea for eigenproblems. The method showed strange behaviour, because of rounding errors, and in an influential paper by Engeli *et al* in 1959 [37], it was shown by careful experiments that the theoretical finite termination within n steps, had no practical meaning. For a discretized biharmonic problem of order 64, they observed that hundreds of steps where necessary in order to obtain the 64 eigenvalues (together with extraneous other values). Wilkinson also analyzed the method and showed that it was forward unstable, which seemed to mark more or less the end of the Lanczos method. It was Paige, who showed in 1971 [97] that the Lanczos method could be used in a truly iterative way in order to obtain correct eigenvalue information. The crux of his analysis is that the observed loss of orthogonality in the Lanczos process, the source of all problems in the method, marked the convergence of an eigenpair; and most remarkably, it did not prevent convergence of other eigenpairs. This loss of orthogonality, by the re-introduction of components of the converged eigenvector to the process, led to duplication of the converged eigenpair in the reduced tridiagonal matrix. The main effect on the convergence of the left eigenpairs seemed to be some delay in the process in exact computation. Paige's analysis spurred much activity in this field and eventually the Lanczos method became a powerful tool and the method of choice for large sparse symmetric matrices, from 1980 on.

We mention the following major steps that led to improvements in the method, and to a better understanding. Parlett and Scott [103] proposed removing the delaying effects on convergence by a *selective* orthogonalization process for the Lanczos vectors. Also, the determination of converged eigenvalues became easier to identify by efficiently computing upper bounds for the residual of an eigenpair. Kaniel (1966) [69] derived upper bounds for the error in an eigenvalue approximation (the so-called *Ritz values*, the eigenvalues of the reduced tridiagonal matrix). These upper bounds have no direct practical implication, since they are in terms of unknown quantities associated with the matrix, such as gaps between eigenvalues relative to the span of the eigenvalues. However, these upper bounds are, in a general sense, rather sharp and can be used for the study of the convergence behaviour of the Ritz values. Later, Saad (1980) [114] refined these upper bounds. The convergence behaviour of Ritz values can be quite irregular; a temporary (almost) stagnation of the process can take place (also called *misconvergence*). Many of these effects were studied carefully, and explained, in a paper by Van der Sluis and Van der Vorst [142], through a rather complicated model for the Ritz values. This model, however, seemed to be necessary in order to show all the intricate dependencies in the apparently simple Lanczos process. In their paper, is is also shown that the convergence behaviour of the Ritz values is superlinear. No matter how irregular the rate of convergence may seem, on average it becomes faster as the iteration proceeds. Parlett also studied the so-called (occasional) misconvergence of the Ritz values [101]. His model is simpler, but does not explain all effects, such as the possible length of the stagnation phase. An algorithm ready for implementation was published by Parlett and Reid [102]; this algorithm computes upper bounds for the errors in the Ritz values, and exploits the fact that we have to compute these for successive tridiagonal matrices. We are not aware of a publicly available implementation, although we know that the process has been implemented for local use.

Note that, strictly mathematically speaking, the Lanczos process is finite and thus it is not correct to use a notion such as "convergence" or even "superlinear convergence". However, in finite precision and for large values of the order of the matrix, the Ritz values will become close to an eigenvalue and for practical purposes the method behaves like a truly convergent process.

In our discussion of the Power Method, we showed how the use of a set of starting vectors arose quite naturally. This is also possible for the Lanczos (and Arnoldi) method, and this approach leads to block and banded Lanczos (and Arnoldi) methods [47, 141, 108].

For unsymmetric matrices it took longer for similar methods to gain popularity. An influential paper, that helped to promote Arnoldi's method as a useful tool, was published by Saad [115]. The Arnoldi method, for orthogonal reduction to upper Hessenberg form, was not only too expensive if one wanted to know only a few eigenpairs, it also suffered from poor convergence for specific eigenvalue distributions. Well-known is the Saad-Schultz example [117], which is a permuted identity matrix. The method leads to trivial approximations after the first $n-1$ steps, and after n steps all eigenpairs suddenly appear. This however, is at a much higher cost than for Householder's reduction. For this reason, the unsymmetric Lanczos process, also referred to as the *two-sided Lanczos method*, received some attention. Initially, the method was notorious for its break-down possibilities, its behaviour in finite precision arithmetic, and the fact the reduction operators to tridiagonal form are non-orthogonal. Cullum and Willoughby, in 1986 [24], presented a code based on the two-sided Lanczos method, in which they solved a number of practical problems; this included a clever trick for identifying the *spurious* eigenvalues due to rounding errors. The code gained some popularity, for instance for plasma-physics eigenvalue computations [21]. Parlett and co-workers [104] introduced the concept of "look-ahead", mainly in order to improve the numerical stability of the process. The look-ahead strategy, introduced in order to prevent breakdown, was further perfected by Freund and Nachtigal, in 1996 [41]. They published a code based on *quasi-minimization* of residuals, and included look-ahead strategies, in which most of the original Lanczos problems were repaired (but the non-orthogonal reductions were still there). Gutknecht [57] published a thorough theoretical overview of the two-sided Lanczos algorithm and exploited its relation to Padé approximations. This gave a better understanding of look-ahead strategies and the convergence behaviour of the method (in the context of solving linear systems). Block variants of the two-sided Lanczos process were discussed in Day's PhD thesis in 1993; for a further description of the algorithms see [3].

Almost simultaneously, there were efforts to make the Arnoldi method more practical. We mention firstly polynomial preconditioning, discussed extensively in Saad's book [116], which damps unwanted parts of the spectrum, and secondly, sophisticated restarting strategies. The method becomes effective for matrices for which shift-and-invert operations can be applied for given vectors. But the many (increasingly expensive) iterations for relevant problems were a bottleneck.

A real breakthrough for the Arnoldi method was realized by Sorensen [125], in 1991, with the so-called *Implicit Restart Technique*. This is a clever technique by which unwanted information can be filtered away from the process. This leads to a reduced subspace with a basis, for which the matrix still has a Hessenberg form, so that Arnoldi's process can be continued with a subspace (rather than with a single vector as with the more classical restart techniques).

The Arnoldi iteration procedure is often carried out with the shift-and-invert approach. For instance, when solving the generalized eigenproblem $Ax = \lambda Bx$, the method is applied to the operator $(A - \sigma B)^{-1}B$. One step further is the *Cayley* transform $(A - \sigma B)^{-1}(A - \tau B)$, which can be used for emphasizing eigenvalues (near σ) and for damping of eigenvalues (near τ). Both techniques require expensive operations with an inverted operator, but the advantage is much faster convergence. Meerbergen and Roose [89] considered the use of inexact Cayley transforms, realized by a few steps of an iterative method, for Arnoldi's method. This technique bears a close relation to polynomial preconditioning. Ruhe [110] considered a more general shift-and-invert transform, the so-called *Rational Krylov Subspace (RKS) Method*:

$$(\delta_j A - \gamma_j B)^{-1}(\sigma_j A - \rho_j B),$$

in which the coefficients may be different for each iteration step j. It has been shown that by generating a subspace with this operator, the given problem can be reduced to a small projected generalized system

$$(\zeta K_{j,j} - \eta L_{j,j})s = 0,$$

where $K_{j,j}$ and $L_{j,j}$ are upper Hessenberg matrices of dimension j. This small system may be solved by the QZ algorithm in order to obtain approximate values for an eigenpair. The parameters in RKS can be chosen to obtain faster convergence to interior eigenvalues. For a comparison of RKS and Arnoldi, see [110, 109].

In 1975, Davidson, a chemist, suggested an iterative method that had the idea of projection on a subspace in common with the Arnoldi method, but with the subspace chosen differently. Motivated by the observation that the matrices in his relevant applications were (strongly) diagonally dominant, Davidson computed the Ritz pairs (the eigenpairs of the projected matrix), computed the residual $r = (A - \theta I)z$ for a pair of interest (θ, z), and proposed expanding the subspace with the vector $(D_A - \theta I)^{-1}r$ (after proper orthogonalization with respect to the current subspace). The matrix D_A denotes the diagonal of the matrix A. For diagonally dominant matrices this approximates, in some sense, inverse iteration with Rayleigh quotient shifts. The Davidson method [26] became quite popular for certain applications. Although other approximations were suggested (see, e.g., [93],) its convergence behaviour for non-diagonal dominant matrices, or for poor initial starting vectors, was far from guaranteed. It was also puzzling that the "optimal" expansion $(A - \theta I)^{-1}r$, optimal from the inverse iteration point of view, led to the vector z, so that the method stagnated. As a consequence, Davidson's method was not able to find the eigenvalues of a diagonal matrix, a very unsatisfactory situation. The method was not very well understood by numerical analysts and as a consequence we find very little reference to it in the numerical analysis literature. Only after 1990, is there some serious analysis [20], almost simultaneously with a successful improvement of the method.

In 1996, Sleijpen and van der Vorst [121] suggested restricting the expansion of the current subspace to vectors from the space orthogonal to z, which restored a largely forgotten technique used by Jacobi (in 1846). Jacobi took an appropriate unit vector for z, and attempted to find the missing component to make it an eigenvector from the space spanned by the remaining $n-1$ unit vectors. He did this by solving a correction equation for the matrix shifted by θ and then restricted the correction to the subspace orthogonal to the chosen unit vector. This correction equation was solved by Gauss-Jacobi iterations and after each two iterations, Jacobi updated the value for θ. Sleijpen and van der Vorst suggested updating z as well and using the update vector for a subspace. This new Jacobi-Davidson method combined Davidson's idea of taking a different subspace with Jacobi's idea of restricting the search of an update to $z^\perp$. We refer to [121] for details. A variant of the technique, in which the correction is approximately solved with one single invert step with a preconditioner, was suggested by Olsen *et al* in 1990 [94]. An exact solution of the correction equation, or an approximate solution of high accuracy, leads to cubic convergence for a properly selected sequence of θ's if A is symmetric, and to quadratic convergence in the unsymmetric case. In the following years it became clear how to efficiently implement the method with preconditioning [122] and how it could be applied to various other eigenproblems, amongst which are generalized eigenproblems [39] and quadratic eigenproblems [123]. Thus, the transformation of these generalized eigenproblems to standard forms can be avoided. The Jacobi-Davidson method is attractive for large sparse eigenproblems for which shift-and-invert operations are too expensive, and for more unusual eigenproblems. The problem of identifying effective preconditioners for the correction matrix for larger classes of matrices is still largely open.

It is well-known that subspace methods lead to eigenvalue approximations that tend to converge towards exterior eigenvalues and that approximations for interior eigenvalues are difficult to obtain. In principle, it is easy to obtain these by working with A^{-1}, but this may be expensive. It is also possible to obtain eigenvalue approximations that converge (slowly) to the eigenvalues of A closest to the origin, from the subspaces generated by A. We explain this for the Arnoldi process. The Arnoldi process leads to

$$AV_m = V_{m+1}H_{m+1,m},$$

where $H_{m+1,m}$ is an upper Hessenberg matrix with $m+1$ rows and m columns. This means that we have a basis for the space with basis vectors Av_j and, using the above relation, this basis can be easily transformed into an orthogonal basis. This orthogonal basis can be used for the projection of A^{-1}, and multiplication by the inverse can be avoided, since all basis vectors have a factor A in common. The exterior eigenvalues of the projected A^{-1}, that is the inverses of interior eigenvalues of A, converge (slowly) to the exterior eigenvalues of A^{-1}. This way of approximating interior eigenvalues has received some attention in the 1990s. In [98], these eigenvalue approximations, in connection with the related Lanczos process, were called Harmonic Ritz values, and some nice relations for Harmonic Ritz values for symmetric indefinite matrices are given in that paper. Harmonic Ritz values had already been studied from a different viewpoint by other authors. Freund [42] studied them as the zeros of the GMRES and MINRES iteration polynomials. Morgan [92] had observed that the Harmonic Ritz values and vectors are very suitable for restarting purposes if one wants to compute interior eigenvalues with subspaces of restricted dimension. In [121, 120], the Harmonic Ritz values are considered in connection with the Jacobi-Davidson process for the selection of proper shifts.

9 Related Topics

9.1 The Singular Value Decomposition

The Singular Value Decomposition plays an essential role in many situations in which we want to decide between relevant and less relevant information. In addition to traditional applications such as regularization of linear least squares problems, and the determination of the numerical rank of a matrix, there are applications to the reduction of information for images and information retrieval from large data bases. A headline contribution to the SIAM News by Berry and Dongarra, summer 1999, showed that the SVD can even help to reduce work in the organization of a conference. Interesting and unusual applications are also shown in a paper by Moler and Morrison [90]. A modern treatment of the numerical aspects of the SVD can be found in Golub and Van Loan's textbook [51] and we have taken most of our information from that source. Demmel [28] gives a good treatment of the important implementational aspects of the SVD.

Since the square of the singular values and the right singular vectors are the eigenpairs for the matrix A^TA (for the left singular vectors this holds with AA^T), it is not surprising that the numerical treatment of the SVD has many relationships to algorithms for symmetric eigenproblems.

This is most visible in Demmel's book, where the discussion on modern SVD algorithms almost parallels the symmetric eigenvalue problem discussion [28, p.211 and p.241]. We note here that working directly with either $A^T A$ or AA^T is not satisfactory for stability and complexity reasons, and this makes separate treatment of the numerical SVD necessary.

The origins of the singular value decomposition go back to the late nineteenth century, with work of Beltrami in 1873. Stewart [128] gave a historic overview of the SVD. The important numerical developments on the SVD started with work of Golub and Kahan in 1965 [49]. This led first to a contribution to the famous ACM-collection as Algorithm 358 [15] and later to the basis of the EISPACK and LAPACK routines in the Wilkinson-Reinsch collection [158, pp.1334-1351]. The key trick in the numerical computation of the SVD is, instead of tridiagonalizing $A^T A$, to bidiagonalize the matrix A. Then, with the bidiagonal reduced forms obtained by orthogonal transformations, one can make variants of QR, divide and conquer, and bisection techniques. The choice between these techniques can be made as above for the symmetric eigenproblem. In particular, for matrices of order larger than 25, the divide and conquer approach is currently regarded as the fastest option [28, p.241].

9.2 Nonlinear eigenproblems and related problems

Standard and generalized eigenproblems arise, for instance, in the study of conservative mechanical systems, governed by Lagrange's equations of small free motion. According to Rogers [106], Rayleigh (in 1873) could not apply his technique for non-conservative systems (systems with a damping term). The well-known technique for the numerical solution of the resulting quadratic eigenproblem is to rewrite it as a generalized eigenproblem [43, 77]. That is

$$\lambda^2 Mx + \lambda Cx + Kx = 0,$$

is equivalent to

$$Az = \lambda Bz$$

with

$$A = \begin{bmatrix} 0 & I \\ -K & -C \end{bmatrix}, \quad B = \begin{bmatrix} I & 0 \\ 0 & M \end{bmatrix}, \quad z = \begin{bmatrix} x \\ \lambda x \end{bmatrix}.$$

Of course, algorithms and theory for the generalized eigenproblem can be used directly, but the unsatisfactory aspect of this is that we have to double the dimension of our spaces. Also, if the generalized eigenproblem is solved approximately by some iteration technique, then it is not straight forward to reduce the information from the double dimensioned space to the original space. Duffin [34] seemed to be the first to generalize the Rayleigh-Ritz principle for the quadratic eigenproblem for symmetric K, C, and M, and A positive definite. Since then, the quadratic eigenproblem, and higher order polynomial problems, have received attention in the numerical literature. Rogers [106], in 1964, considered a more general quadratic eigenproblem and used minimax principles for the investigation of such problems. In the late sixties and early seventies, algorithms that avoid the linearization step appeared. We mention the work of Lancaster [77], Kublanovskaja [75, 76], and Ruhe [107]. The suggested algorithms are mostly variants of Newton's method. More recently, in 1995, Guo *et al* [56] described several iterative methods, that can be regarded as a fixed point iteration combined with the Lanczos method and a (simplified) Newton iteration. A backward error analysis for more general polynomial eigenproblems was given by Tisseur (1998 [133]), and a perturbation analysis for quadratic problems was published by Sun in 1999 [130]. In a paper that appeared in 1996, Sleijpen *et al* [120] showed that the Jacobi-Davidson method could be applied in order to reduce a given polynomial problem in an n-dimensional space to a similar problem in a much lower m-dimensional space. The problem in the lower dimensional space can then be attacked by any of the previously mentioned approaches. For a quadratic equation from an accoustic problem, it was shown how this approach led successfully and efficiently to the desired eigenpair, for matrices of the order of about $240,000$ [123]. For higher polynomial eigenproblems there has not been much experience to date. Bai [6] mentioned the need for algorithms for fifth order polynomial eigenproblems in a review paper that appeared in 1995. Quite recently, in 1998, Heeg [58] showed how the Jacobi-Davidson approach could be successfully applied to fourth-order polynomial eigenproblems with complex matrices in the study of instabilities of attachment-line flows for airfoils.

Other than these polynomial eigenproblems, there is a wide variety of problems that are associated with the determination of invariant subspaces. Since the standard and generalized eigenproblems became more familiar and more or less routinely solvable, the more difficult problems received more attention. Among these problems we mention the following; the *Procrustes problem*: minimize $\|AY - YB\|_F$ for given A and B over the manifold $Y^*Y = I$; the determination of a nearest Jordan structure, and the problem of determining a simultaneous Schur decomposition for a set of perturbed matrices (under the assumption that the unperturbed matrices have a simultaneous Schur form). A nice overview of such problems, as well as software for the numerical solution is described by Edelman and Lippert [83].

9.3 Pseudospectra

Eigenvalues are often used as a source of information on stability or convergence, and the question arises as to the validity of the information gained from these values. For example, during the 1990s it was realized (see, e.g., [138]) that eigenvalues alone do not govern the instability and transition to turbulence of high Reynolds number fluid flows as had previously been thought. Many authors have studied the problem of sensitivity of the eigenvalues with respect to perturbations, see for instance [154, 129, 19]. These studies are usually related to perturbations caused by rounding errors, and not so much by the relevance of the eigenvalues due to the particular representation of a given problem, for instance the choice of basis.

Around 1987, Trefethen [140] began to emphasize this aspect of eigencomputations. He propagated the idea of inspecting the pseudospectra of a matrix as a relatively simple means for determining the significance of a particular part of the spectrum, without getting involved in complicated matters such as angles between eigenvectors or eigenspaces[5]. The definition of the pseudospectrum $\Lambda_\varepsilon(A)$ for a matrix A is directly related to perturbations:

$$\Lambda_\varepsilon(A) \equiv \{z \in \mathbb{C} : z \in \Lambda(A + E) \text{ for some } E \text{ with } \|E\| \le \varepsilon\}.$$

The pseudospectra are usually shown graphically as a set of level curves for various values of ε. The level curves, or contour integrals, are more apparent from the original definition of the ε-pseudospectrum, in terms of the norm of the resolvent $(zI - A)^{-1}$:

$$\Lambda_\varepsilon(A) \equiv \{z \in \mathbb{C} : \|(zI - A)^{-1}\| \ge \varepsilon^{-1}\},$$

with the convention $\|(zI - A)^{-1}\| = \infty$ for $z \in \Lambda(A)$.

For symmetric matrices, the pseudospectra of A is a collection of discs around the eigenvalues of A (note that the perturbation E need not be symmetric). For unsymmetric matrices the pseudospectra can be any collection of curves around the set of eigenvalues of A. These level curves may give information that is hidden by the information provided by the eigenvalues themselves. For instance, when studying stability of integration methods for systems of ODE's, or in bifurcation problems, the eigenvalues may be in a proper region, for instance, in the left-half plane, while the level curves, even for small values of ε, may intersect with the right-half plane. This suggests that it may be necessary to ask further questions about the problem. On the other hand, the pseudospectra may not tell the full story. For instance, the sensitivity problems may be due to a single pair of ill-conditioned eigenvectors for which the more global level curves are too pessimistic. It may be the case that it is not realistic to assume equal perturbations for all matrix entries, but nevertheless the pseudospectra focus attention on critical places in the spectrum. A nice introduction to the

[5] The notion had been defined by others earlier, the earliest of them, according to Trefethen, being Landau [79]

relevance of pseudospectra is given in [135], where the pseudospectra are actually computed and discussed for a number of matrices.

Due to the nature of computing pseudospectra, this useful tool is often restricted to matrices of relatively moderate size and one has to be careful in generalizing the insights gained from smaller problems to larger problems which are similar. More recently, tools have become available for computing the pseudospectrum of large sparse matrices. Carpraux *et al* [16] proposed an algorithm for computing the smallest singular value of $zI - A$, that is based on Davidson's method using ILU preconditioning. Lui [84] (see also [13]) suggested using the Lanczos method in combination with continuation techniques. This is a plausible approach, since we need to do the computation for many values of z, well-distributed over the region of interest, in order to obtain a complete picture of the pseudospectra. We know that currently such tools are being used for the analysis of instability problems of large sets of ODE's, related to climate modelling, but results have not yet been published. Valuable information on pseudospectra can also be derived from the Arnoldi iteration obtained, for instance, with ARPACK [81]. For an up-to-date overview on pseudospectra results, see Trefethen's 1999 paper [136].

9.4 Homotopy methods

The subspace methods that we have discussed above are often applied in combination with shift-and-invert operations. This means that if one wants to compute eigenvalues close to a value σ, then the methods are applied to $(A - \sigma I)^{-1}$. As we have seen, the Jacobi-Davidson method can be interpreted as an inexact shift-and-invert method, since the invert step is usually approximated by a few steps of some convenient preconditioned inner iteration method.

Another method related to these inexact shift-and-invert approaches is the homotopy approach which has received some attention in the nineties. The idea is to compute some of the eigenvalues of a perturbed matrix $A + E$, when the eigenvalues of A are known, or can be relatively easily computed. In order to this we use the homotopy $H(t) = A + tE$, $0 \leq t \leq 1$. If eigenpairs of $H(t_0)$ are known, then they are used as approximations for those of $H(t_0 + \Delta t)$. These approximations are improved by a convenient subspace iteration (cf. [85]). Rayleigh quotient iterations are used for symmetric A and E (see references in [85] for earlier work on homotopy for eigenproblems). For the Rayleigh quotient iteration one needs to solve systems like $(H(t_0+\Delta t) - \lambda I)y = x$, where (λ, x) represents the current approximation for an eigenpair of $H(t_0 + \Delta t)$. In the context of large sparse matrices, it may be undesirable to do this with a direct solver, and in [85] the system is solved with SYMMLQ [99]. Of course, one could restrict oneself to only a few steps with SYMMLQ, and then try to accelerate the inexact Rayleigh quotient steps, as is done in the Jacobi-Davidson method. This indicates relations between these different approaches, but as far as we know, these relations have not yet been explored. In [85], it is observed that SYMMLQ may have difficulty in converging for the nearly singular system $H(t_0 + \Delta t) - \lambda I)y = x$, and it is suggested that the situation would be improved by applying the Rayleigh quotient iteration to the approximately deflated matrix $H(t_0 + \Delta t) + xx^T$. (The term approximately deflated is used to indicate that x is only an approximation to the desired eigenvector.) Note that similar deflation procedures are incorporated in the Jacobi-Davidson process. The whole procedure is repeated for successive increments Δt, until the final value $t = 1$ is reached. In [85] an elegant approach is followed for the selection of the step size Δt.

The homotopy approach lends itself quite naturally to situations where the matrix A varies in time, or where it varies as a linearization of a nonlinear operator, as in bifurcation problems. Another example of an interesting problem is the Schrödinger eigenvalue problem [85]:

$$-\Delta u + fu = \lambda u$$

in the unit square in two dimensions with homogeneous Dirichlet boundary conditions. With the usual finite difference approximations on a uniform grid, this leads to the discrete Laplacian for $-\Delta u$, for which we know the eigensystem.

In [159], the homotopy approach is used for symmetric generalized eigenproblems, very much along the lines sketched above. The application to real unsymmetric eigenproblems is considered in [86].

In [82], the homotopy approach is suggested as a means of realizing a divide and conquer method for unsymmetric eigenproblems, as an alternative for the solution of the secular equation for symmetric problems (which cannot be used for unsymmetric problems).

9.5 Miscellaneous

Subspace methods such as Simultaneous Iteration, RQI, Lanczos, Arnoldi, and Davidson for large sparse systems are in general more powerful in combination with shift-and-invert techniques. This requires the solution of large sparse linear systems, since these methods only need operations with, for instance, $(A - \sigma I)$ in order to compute matrix vector products $(A - \sigma I)^{-1}v$. This has led to the question of how accurately these operations have to be carried out in order to maintain a reasonable rate of convergence. For instance, for the Power Method, this approximation technique, or preconditioning, can be described as follows for the standard eigenproblem $Ax = \lambda x$. For each iteration one computes, for the current iteration vector $x^{(i)}$ the Rayleigh quotient $\mu^{(i)} = \frac{x^{(i)*}Ax^{(i)}}{x^{(i)*}x^{(i)}}$. Instead of computing $x^{(i+1)}$ with $Ax^{(i)}$, the defect $r = Ax^{(i)} - \mu^{(i)}x^{(i)}$ is computed and this defect is multiplied by the preconditioner K. The new vector $x^{(i+1)}$ is then computed as a normalized linear combination of x^i and Kr. Note the resemblance to Davidson's approach.

The convergence of this basic algorithm, including its use for the generalized eigenproblem, has been studied quite extensively in the Russian literature, starting in about 1980 in the work of D'yakonov [35] and others. Knyazev [72] gave an excellent overview of the work in the Russian literature on *preconditioned eigensolvers*. The study of these preconditioned iteration techniques is relevant also to the understanding of inexact, or preconditioned, forms of shift-and-invert Lanczos (proposed by Morgan and Scott [93]), inexact Arnoldi (work by Meerbergen [87]) and Davidson variants, including the Jacobi-Davidson method. A presentation of various iteration techniques including preconditioning was given by Knyazev in [71]. Smit [124, Ch.4] studied the effect of approximate inverses on the convergence of the RQI method. For a discussion on inexact Krylov methods, see Meerbergen and Morgan [88]. Preconditioning in relation to the Jacobi-Davidson method is discussed in [122]. Note that in all these algorithms, the preconditioner is used for the computation of a promising new direction vector; the given eigenproblem is untouched. This is different from the situation with linear solution methods, where preconditioning is used to transform the given system to one that can be handled more efficiently.

10 Software

The history of reliable high quality software for numerical linear algebra started with the book edited by Wilkinson and Reinsch, the *Handbook for Automatic Computation, Vol 2, Linear Algebra*, published in 1971. This book contained a number of articles that had appeared previously in Numerische Mathematik, which described state-of-the-art algorithms for the solution of linear systems and eigenproblems. All these articles contained implementations in Algol60. Most of these algorithms are still alive albeit in other languages. Algol60 was a computer language that gained some popularity in academia, mostly in Europe, but it was not as fast as Fortran on most machines and it did not gain a foothold in the slowly emerging large industrial codes (the majority of which were written in Fortran or even in assembler language). For this reason, groups in the USA started the development of two influential software packages LINPACK and EISPACK in the early 1970s. These packages started as transcriptions of the major part of the Wilkinson and Reinsch collection: LINPACK covered the numerical solution of linear systems; EISPACK concentrated on eigenvalue problems. The most prominent omissions form these packages were iterative solution methods: the conjugate gradient method was not included in LINPACK and the Jacobi method was not included in EISPACK. At that time the Lanczos and Arnoldi methods were not even considered as candidates, because they were viewed either as direct methods, and in that sense not competitive with the then available methods, or as iterative methods that could not be safely automated [68]. The Lanczos method was only considered, in that time, as a safe iterative method providing one did complete re-orthogonalization (cf. [48]).

The "Wilkinson and Reinsch" procedures can also be viewed as prototypes for eigenvalue routines in the bigger software packages NAG and IMSL, and in the widely available software package MATLAB. EISPACK was replaced in 1995 by LAPACK, in the words of Golub and Van Loan (1996): *"LAPACK stands on the shoulders of two other packages (viz: LINPACK and EISPACK) that are milestones in the history of software development"*. A more recent development along these lines is ScaLAPACK [11] which aims to provide close to optimal software for modern parallel computers.

We have already mentioned MATLAB in passing; the impact of this computing environment in the scientific computing world has been tremendous. MATLAB provides state-of-the-art software for all sorts of numerical linear algebra computations and has become the *de facto* standard for coding and testing in the 1990s. Its impact on the development of this field can hardly be overestimated.

As already indicated, in 1970 there were few robust and well-understood iterative methods available, and mainly for that reason, these methods were not included in the packages constructed then. They are still not available in their successors, with the exception of MATLAB (since 1998, MATLAB has had iterative methods for eigenproblems available). It was soon clear that the powerful "direct" methods, based on reduction to some special form, had their limitations for the large sparse matrices that occur in the modeling of realistic stability problems, and there was a heavy demand for methods that could, at least partially, handle big problems. Cullum and Willoughby filled this gap, in 1985, with their software based on the two-sided Lanczos procedure [22, 23, 24]. Of course, they realized the intricacies of the Lanczos method, and they advocated the use of iterated Rayleigh quotient steps for improvement of the information of the Lanczos output. They also provided software for this updating step (their software was available some years before 1985, but its publication in book form took place in that year).

Freund and Nachtigal proposed in 1992, a variant of the two-sided Lanczos process that improved convergence properties using a quasi minimization step; they also included sophisticated look-ahead facilities. The QMR method could also be used for eigenvalue computations, and they provided software through Netlib for this purpose, see QMRPACK [41]. Algorithms and software for adaptive block-variants of the two-sided Lanczos algorithm (ABLE) have been described Bai, Day and Ye [3].

Because of the improvements made to the Arnoldi method, in particular the implicit restart technique, it became feasible to exploit the attractive orthogonal reduction properties in an efficient manner. This was realized in the ARPACK software, for which the User's guide was published, by Lehoucq, Sorensen, and Yang, in 1998 [81]. The package was in existence and available to users a few years earlier. At present, ARPACK seems to be the default choice for large sparse eigenproblems, provided that either it is possible to implement shift-and-invert operations efficiently, or that unwanted parts of the spectrum can be damped by a fixed (polynomial) preconditioner. The parallel implementation of ARPACK is referred to as P_ARPACK, it is portable across a wide range of distributed memory platforms. The parallelism is mainly in the matrix vector products and the user has full control over this trough the reverse communication principle. For more details, see [81, 31].

Kågström *et al* have provided software for the generalized upper staircase Schur form. The software, called GUPTRI, can be used for the analysis of singular pencils $A - \lambda B$. This reduction admits more general cases than the QZ decomposition, for instance A and B may be singular. This software has been described in [4]. See also this reference for pointers to recent work by Kågström.

For a thorough overview of modern algorithms, software and pointers to sources for further information see [4].

11 Epilogue

For symmetric matrices, the eigenproblem is relatively simple, due to the existence of a complete orthogonal eigensystem, and the fact that all eigenvalues are real. These properties are exploited in the most efficient numerical methods, and the symmetric eigenproblem may be considered as solved: for small matrices $n \leq 25$ we have the QR method, one of the most elegant numerical techniques produced in the field of numerical analysis; for larger matrices (but smaller than a few thousand), we have a combination of divide and conquer with QR techniques. For the largest matrices, there is the Lanczos method, which in its pure form is strikingly simple but which conceals so many nice and attractive properties. All methods have a relatively long history. The theoretical aspects of the computations (convergence and perturbation theory) are relatively well understood.

For unsymmetric matrices the picture is less rosy. Unfortunately, it is not always possible to diagonalize an unsymmetric matrix, and even if it is known that all eigenvalues are distinct, then it may be numerically undesirable to to this. The most stable methods seem to be based on the Schur factorization, that is for each n by n matrix A, there exists an orthogonal Q, so that

$$Q^* A Q = R,$$

in which R is upper triangular. Apart from permutations and signs, the matrix Q is unique. The Schur factorization reveals much of the eigenstructure of A: its diagonal elements are the eigenvalues of A and the off-diagonal elements of R indicate how small the angles between eigenvectors may be. For matrices not too large, QR is the method of choice, but for larger matrices the picture is less clear. Modern variants of the Arnoldi method seem to be the first choice at the moment, and, if approximations are available, the Jacobi-Davidson method may be attractive. There is still a lot to investigate: if the matrix is nonnormal, that is, if the eigenvectors do not form a unitary set, then what kind of meaningful information can we extract from a given matrix (invariant subspaces, angles between subspaces, distance to Jordan canonical form), and how can we compute this as accurately as possible? Much has been done, but even more remains to be done.

There are even more open problems as eigenproblems become more complicated: generalized eigenproblems, polynomial eigenproblems, nonlinear eigenproblems, etc.

Looking back over the past century, we see that the solution of the eigenproblem has given up some of its mysteries through the work of many devoted and highly talented researchers. Novices in the field should be aware that the modern algorithms, even the apparantly simple ones, are the result of many independent "small" steps. The fact that many of these steps can now be regarded as "small" illustrates how theory has kept up the pace with computational practice, so that new developments can find their place in an expanding but still elegant framework. Astonishingly much has been achieved, both computationally and theoretically, in a concerted effort, but much more remains to be unravelled. It is our firm belief that eigenproblem research will remain a lively and useful area of research for a long time to come. We hope that this overview will help to motivate young researchers to make their contributions to solving pieces of the gigantic puzzles that remain.

Acknowledgements

We wish to thank our colleagues Michele Benzi, Jan Brandts, Mark Embree, Karl Meerbergen, Gerard Sleijpen, Nick Trefethen, and David Watkins, for their suggestions for improvements and corrections.

References

[1] W. E. Arnoldi. The principle of minimized iteration in the solution of the matrix eigenproblem. *Quart. Appl. Math.*, 9:17–29, 1951.

[2] E. Bodewig. *Matrix Calculus*. North-Holland, Amsterdam, 1956.

[3] Z. Bai, D. Day, and Q. Ye. ABLE: An adaptive block Lanczos method for non-Hermitian eigenvalue problems. Technical Report Research Report 95-04, University of Kentucky, Lexington, KY, 1995.

[4] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst. *Templates for the Solution of Algebraic Eigenvalue Problems: a Practical Guide*. SIAM, Philadelphia, PA, 1999.

[5] Z. Bai and G.W. Stewart. SRRIT - A FORTRAN subroutine to calculate the dominant invariant subspace of a nonsymmetric matrix. Technical Report 2908, Department of Computer Science, University of Maryland, 1992.

[6] Zhaojun Bai. Progress in the numerical solution of the nonsymmetric eigenvalue problem. *Num. Lin. Alg. Appl.*, 2:219–234, 1995.

[7] V. Bargmann, C. Montgomery, and J. von Neumann. Solution of linear systems of high order. Technical report, Princeton, Institute for Advanced Study, 1946.

[8] F. L. Bauer. Das Verfahren der Treppeniteration und verwandte Verfahren zur Lösung algebraischers Eigenwertprobleme. *Z. angewandte Math. Phys.*, 8:214–235, 1957.

[9] F. L. Bauer and C. T. Fike. Norms and exclusion theorems. *Numer. Math.*, 2:137–141, 1960.

[10] C. H. Bischof. The two-sided block Jacobi method on hypercube architectures. In M. T. Heath, editor, *Hypercube Multiprocessors*, Philadelphia, PA, 1987. SIAM.

[11] L. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. Whaley. *ScaLAPACK User's Guide*. SIAM, Philadelphia, PA, 1997.

[12] T. Braconnier, F. Chatelin, and V. Frayssé. The influence of large nonnormality on the quality of convergence of iterative methods in linear algebra. Technical Report TR-PA-94/07, CERFACS, Toulouse, 1994.

[13] T. Braconnier and N. J. Higham. Computing the field of values and pseudspectra using the Lanczos method with continuation. *BIT*, 36:422–440, 1996.

[14] R. P. Brent and F. T. Luk. The solution of singular value and symmetric eigenvalue problems on multiprocessor arrays. *SIAM J. Sci. Statist. Comput.*, 6:69–84, 1985.

[15] P. A. Businger and G. H. Golub. Algorithm 358: Singular value decomposition of a complex matrix. *Comm. Assoc. Comp. Mach.*, 12:564–565, 1969.

[16] J. F. Carpraux, J. Erhel, and M. Sadkane. Spectral portrait for non-Hermitian large sparse matrices. *Computing*, 53:301–310, 1994.

[17] F. Chaitin-Chatelin. Is nonnormality a serious difficulty? Technical Report TR-PA-94/18, CERFACS, Toulouse, 1994.

[18] F. Chatelin. *Spectral Approximation of Linear Operators*. Academic Press, New York, 1983.

[19] F. Chatelin. *Eigenvalues of Matrices*. J. Wiley and Sons, Chichester, 1993. Masson, 1988.

[20] M. Crouzeix, B. Philippe, and M. Sadkane. The Davidson method. *SIAM J. Sci. Comput.*, 15:62–76, 1994.

[21] J. Cullum, W. Kerner, and R. Willoughby. A generalized nonsymmetric Lanczos procedure. *Computer Physics Communications*, 53:19–48, 1989.

[22] J. Cullum and R. A. Willoughby. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Vol I*. Birkhaüser, Boston, MA, 1985.

[23] J. Cullum and R. A. Willoughby. *Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Vol II*. Birkhaüser, Boston, MA, 1985.

[24] J. Cullum and R. A. Willoughby. A practical procedure for computing eigenvalues of large sparse nonsymmetric matrices. In J. Cullum and R. A. Willoughby, editors, *Large Scale Eigenvalue Problems*, pages 193–240, Amsterdam, 1986. North-Holland.

[25] J. J. M. Cuppen. A divide and conquer method for the symmetric eigenproblem. *Numer. Math.*, 36:177–195, 1981.

[26] E. R. Davidson. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real symmetric matrices. *J. Comp. Phys.*, 17:87–94, 1975.

[27] C. Davis and W. Kahan. The rotation of eigenvectors by a perturbation-iii. *Soc.Ind.Appl.Math.J. Num.Anal.*, 7:1–46, 1970.

[28] J. W. Demmel. *Applied Numerical Linear Algebra.* SIAM, Philadelphia, PA, 1997.

[29] J. W. Demmel and K. Veselić. Jacobi's method is more accurate than QR. *SIAM J. Matrix Anal. Appl.*, 13:1204–1246, 1992.

[30] A. den Boer. *De Jacobi-methode van 1845 tot 1990.* Master Thesis, Dept. of Mathematics, Utrecht University, 1991.

[31] J. J. Dongarra, I. S. Duff, D. C. Sorensen, and H. A. van der Vorst. *Numerical Linear Algebra for High-Performance Computers.* SIAM, Philadelphia, PA, 1998.

[32] J. J. Dongarra and D. C. Sorensen. A fully parallel algorithm for the symmetric eigenvalue problem. *SIAM J. Sci. Statist. Comput.*, 8:139–154, 1987.

[33] J. J. Dongarra and D. W. Walker. Software libraries for linear algebra computations on high performance computers. *SIAM Rev.*, 37:151–180, 1995.

[34] R. J. Duffin. A minimax theory for overdamped networks. *J. of Rational Mechanics and Analysis*, 4:221–233, 1955.

[35] E. G. D'yakonov. Iterative methods in eigenvalue problems. *Math. Notes*, 34:945–953, 1983.

[36] P. J. Eberlein. Solution to the complex eigenproblem by a norm-reducing Jacobi-type method. *Numer. Math.*, 14:232–245, 1970.

[37] M. Engeli, T. Ginsburg, H. Rutishauser, and E. Stiefel. *Refined Iterative Methods for Computation of the Solution and the Eigenvalues of Self-Adjoint Boundary Value Problems.* Birkhäuser, Basel/Stuttgart, 1959.

[38] D. G. Feingold and R. S. Varga. Block diagonally dominant matrices in generalizations of the Gerschgorin circle theorem. *Pacific J. Math.*, 12:1241–1250, 1962.

[39] D. R. Fokkema, G. L. G. Sleijpen, and H. A. van der Vorst. Jacobi-Davidson style QR and QZ algorithms for the reduction of matrix pencils. *SIAM J. Sci. Comput.*, 20:94–125, 1996.

[40] J.G.F. Francis. The QR transformation: a unitary analogue to the LR transformation, Parts I and II. *Comp. J.*, 4:265–272, 332–345, 1961.

[41] R. W. Freund and N. M. Nachtigal. QMRPACK: a package of QMR algorithms. *ACM Trans. Math. Softw.*, 22:46–77, 1996.

[42] R.W. Freund. Quasi-kernel polynomials and their use in non-Hermitian matrix iterations. *J. Comp. and Appl. Math.*, 43:135–158, 1992.

[43] F. Gantmacher. *The theory of matrices, vol I and II.* Chelsea, New York, 1959.

[44] S. Gerschgorin. Über die Abgrenzung der Eigenwerte einer Matrix. *Izv. Akad. Nauk SSSR, Ser. fiz.-mat.*, 6:749–754, 1931.

[45] W. Givens. Numerical computation of the characteristic values of a real symmetric matrix. Oak Ridge National Lab., ORNL-1574, 1954.

[46] H. H. Goldstine, F. J. Murray and J. von Neumann. The Jacobi method for real symmetric matrices. *J. Ass. Comp. Mach.*, 6:176–195, 1959.

[47] G. H. Golub. Some uses of the Lanczos algorithm in numerical linear algebra. In J. J. H. Miller, editor, *Topics in Numerical Analysis*, pages 173–184. Academic Press, New York, 1973.

[48] G. H. Golub. Sparse matrix computations: Eigenvalues and linear equations. In *Seminaries Iria*, Inst. Rech. D'Informatique et d'Automatique, Rocquencourt, pages 117–140, 1975.

[49] G. H. Golub and W. Kahan. Calculating the singular values and pseudo-inverse of a matrix. *SIAM J. Numer. Anal.*, 2:205–224, 1965.

[50] G. H. Golub and H. A. van der Vorst. 150 Years old and still alive: Eigenproblems. In *The State of the Art in Numerical Analysis*, pages 93–119. Clarendon Press, Oxford, 1997.

[51] G. H. Golub and C. F. Van Loan. *Matrix Computations.* The Johns Hopkins University Press, Baltimore, 1996.

[52] G. H. Golub and J. H. Wilkinson. Ill-conditioned eigensystems and the computation of the Jordan canonical form. *SIAM Rev.*, 18:578–619, 1976.

[53] J. Greenstadt. A method for finding roots of arbitrary matrices. *MTAC*, 9:47–52, 1955.

[54] R. T. Gregory. Computing eigenvalues and eigenvectors of a symmetric matrix on the ILLIAC. *MTAC*, 7:215–220, 1953.

[55] M. Gu and S. C. Eisenstat. A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem. *SIAM J. Matrix Anal. Applic.*, 16:172–191, 1995.

[56] J.-S. Guo, W.-W. Lin, and C.-S. Wang. Numerical solutions for large sparse eigenvalue problems. *Linear Alg. Appl.*, 225:57–89, 1995.

[57] M. H. Gutknecht. A completed theory of the unsymmetric Lanczos process and related algorithms, Part I. *SIAM J. Matrix Anal. Applic.*, 13:594–639, 1992.

[58] R. Heeg. *Stability and transition of attachment-line flow.* PhD thesis, Universiteit Twente, Enschede, The Netherlands, 1998.

[59] P. Henrici. On the speed of convergence of cyclic and quasicyclic Jacobi methods for computing the eigenvalues of Hermitian matrices. *SIAM J. Appl. Math.*, 6:144–162, 1958.

[60] K. Hessenberg. *Auflösung linearer Eigenwertaufgaben met Hilfe der Hamilton-Cayleyschen Gleichung.* PhD Thesis, Technische Hochschule, Darmstadt, 1940.

[61] W. Hoffmann and B. N. Parlett. A new proof of global convergence for the tridiagonal QL algorithm. *SIAM J. Numer. Anal.*, 15:929–937, 1978.

[62] P. Horst. A method of factor analysis by means of which all coordinates of the factor matrix are given simultaneously. *Psychometrika*, 2:225–236, 1937.

[63] A. S. Householder. Unitary triangularization of a nonsymmetric matrix. *J. ACM*, 5:339–342, 1958.

[64] A. S. Householder. *The Theory of Matrices in Numerical Analysis.* Dover, New York, 1964.

[65] C. G. J. Jacobi. Über ein leichtes Verfahren, die in der Theorie der Säcularstörungen vorkommenden Gleichungen numerisch aufzulösen. *Journal für die reine und angewandte Mathematik*, pages 51–94, 1846.

[66] B. Kågström. The generalized singular value decomposition and the general $a - \lambda b$ problem. *BIT*, 24:568–583, 1985.

[67] B. Kågström. RGSVD: An algorithm for computing the Kronecker structure and reducing subspaces of singular $a - \lambda b$ pencils. *SIAM J. Sci. Statist. Comput.*, 7:185–211, 1986.

[68] W. Kahan and B. N. Parlett, How far should you go with the Lanczos process. In J. R. Bunch and D. J. Rose, editors, *Sparse Matrix Computations.* Academic Press, New York, 1976.

[69] S. Kaniel. Estimates for some computational techniques in linear algebra. *Math. Comp.*, 20:369–378, 1966.

[70] T. Kato. On the upper and lower bounds of eigenvalues. *J. Phys. Soc. Japan*, 4:334–339, 1949.

[71] A. Knyazev. Preconditioned eigensolvers. In Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, editors, *Templates for the Solution of Algebraic Eigenvalue Problems: a Practical Guide.* SIAM, Philadelphia, PA, to appear in 2000.

[72] A. V. Knyazev. Preconditioned eigensolvers - an oxymoron? *ETNA*, 7:104–123, 1998.

[73] A. N. Krylov. O cislennom resenii uravnenija, kotorym v techniceskih voprasah opredeljajutsja castoty malyh kolebanii material'nyh sistem. *Izv. Akad. Nauk SSSR. Ser. fiz.-mat.*, 4:491–539.

[74] V. N. Kublanovskaja. On some algorithms for the solution of the complete eigenvalue problem. *Zh. Vych. Mat.*, 1:555-570, 1961.

[75] V. N. Kublanovskaja. On an application of Newton's method to the determination of eigenvalues of λ-matrices. *Soviet Math. Dokl.*, 10:1240–1241, 1969.

[76] V. N. Kublanovskaja. On an approach to the solution of the generalized latent value problem for λ-matrices. *SIAM J. Numer. Anal.*, 7:532–537, 1970.

[77] P. Lancaster. *Lambda-matrices and vibrating systems.* Pergamon Press, Oxford, UK, 1969.

[78] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Natl. Bur. Stand*, 45:225-280, 1950.

[79] H. J. Landau. On Szegö's eigenvalue distribution theory and non-Hermitian kernels. *J. d'Analyse Math.*, 28:335-357, 1975.

[80] S. L. Lee. A practical upper bound for departure from normality. *SIAM J. Matrix Anal. Applic.*, 16:462–468, 1995.

[81] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK User's Guide - Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods.* SIAM, Philadelphia, PA, 1998.

[82] T.-Y. Li and Z. Zeng. Homotopy-determinant algorithm for solving nonsymmetric eigenvalue problems. *Math. Comp.*, 59:483–502, 1992.

[83] R. Lippert and A. Edelman. Nonlinear eigenvalue problems with orthogonality constraints. In Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, editors, *Templates for the Solution of Algebraic Eigenvalue Problems: a Practical Guide.* SIAM, Philadelphia, to appear.

[84] S. H. Lui. Computation of pseudospectra by continuation. *SIAM J. Sci. Comput.*, 1996.

[85] S. H. Lui and G. H. Golub. Homotopy method for the numerical solution of the eigenvalue problem of self-adjoint partial differential operators. *Numerical Algorithms*, 10:363–378, 1995.

[86] S. H. Lui, H. B. Keller, and W. C. Kwok. Homotopy method for the large sparse real nonsymmetric eigenvalue problem. Technical Report ??, Department of Mathematics, Hong Kong University of Science and Technology, Kowloon, Hong Kong, 1996.

[87] K. Meerbergen. *Robust methods for the calculation of rightmost eigenvalues of nonsymmetric eigenvalue problems.* PhD thesis, Katholieke Universiteit Leuven, Leuven, Belgium, 1996.

[88] K. Meerbergen and R. Morgan. Inexact methods. In Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, editors, *Templates for the Solution of Algebraic Eigenvalue Problems: a Practical Guide.* SIAM, Philadelphia, PA, 1999.

[89] K. Meerbergen and D. Roose. The restarted Arnoldi method applied to iterative linear system solvers for the computation of rightmost eigenvalues. *SIAM J. Matrix Anal. Applic.*, 18:1–20, 1997.

[90] C. B. Moler and D. Morrison. Singular value analysis of cryptograms. *Amer. Math. Monthly*, pages 78–87, 1983.

[91] C. B. Moler and G. W. Stewart. An algorithm for generalized matrix eigenvalue problems. *SIAM J. Numer. Anal.*, 10:241–256, 1973.

[92] R. B. Morgan. Computing interior eigenvalues of large matrices. *Linear Alg. Appl.*, 154–156:289–309, 1991.

[93] R. B. Morgan and D. S. Scott. Generalizations of Davidson's method for computing eigenvalues of sparse symmetric matrices. *SIAM J. Sci. Statist. Comput.*, 7(3):817–825, 1986.

[94] J. Olsen, P. Jørgensen, and J. Simons. Passing the one-billion limit in full configuration-interaction (FCI) calculations. *Chemical Physics Letters*, 169:463–472, 1990.

[95] A. M. Ostrowski. On the convergence of the Rayleigh quotient iteration for the computation of characteristic roots and vectors I–VI. *Arch. Rational. Mesh. Anal.*, I 1:233–241, II 2:423–428, III 3:325–340, IV 3:341–347, V 3:472–481, VI 4:153-165, 1958-1959.

[96] M. H. C. Paardekooper. An eigenvalue algorithm for skew symmetric matrices. *Numer. Math.*, 17:189–202, 1971.

[97] C. C. Paige. *The computation of eigenvalues and eigenvectors of very large matrices.* PhD thesis, Univ. of London, 1971.

[98] C. C. Paige, B. N. Parlett, and H. A. van der Vorst. Approximate solutions and eigenvalue bounds from Krylov subspaces. *Num. Lin. Alg. Appl.*, 2(2):115–134, 1995.

[99] C. C. Paige and M. A. Saunders. Solution of sparse indefinite systems of linear equations. *SIAM J. Numer. Anal.*, 12:617–629, 1975.

[100] B. N. Parlett. *The Symmetric Eigenvalue Problem.* Prentice-Hall, Englewood Cliffs, N.J., 1980.

[101] B. N. Parlett. Misconvergence in the Lanczos algorithm. In M. G. Cox and S. Hammarling, editors, *Reliable Numerical Computation*, Chapter 1, p.7–24. Clarendon Press, Oxford UK, 1990.

[102] B. N. Parlett and J. K. Reid. Tracking the progress of the Lanczos algorithm for large symmetric eigenproblems. *IMA J. Num. Anal.*, 1:135–155, 1981.

[103] B. N. Parlett and D. S. Scott. The Lanczos algorithm with selective orthogonalization. *Math. Comp.*, 33:217–238, 1979.

[104] B. N. Parlett, D. R. Taylor, and Z. A. Liu. A look-ahead Lanczos algorithm for unsymmetric matrices. *Math. Comp.*, 44:105–124, 1985.

[105] D.A. Pope and C. Tompkins. Maximizing functions of rotations: experiments concerning speed of diagonalization of symmetric matrices using Jacobi's method. *J. ACM 4*, 459–466, 1957.

[106] E. H. Rogers. A minimax theory for overdamped systems. *Arch. for Rational Mechanic and Analysis*, 19:89–96, 1964.

[107] A. Ruhe. Algorithms for the nonlinear eigenvalue problem. *SIAM J. Numer. Anal.*, 10:674–689, 1973.

[108] A. Ruhe. Implementation aspects of band Lanczos algorithms for computation of eigenvalues of large sparse symmetric matrices. *Math. Comp.*, 33:680–687, 1979.

[109] A. Ruhe. The rational Krylov algorithm for nonsymmetric eigenvalue problems. iii: Complex shifts for real matrices. *BIT*, 34:165–176, 1994.

[110] A. Ruhe. Rational Krylov algorithms for nonsymmetric eigenvalue problems. ii. Matrix pairs. *Linear Alg. Appl.*, 197/198:283–295, 1994.

[111] H. Rutishauser. Solution of eigenvalue problems with the LR-transformation. *Nat. Bur. Standards Appl. Math. Ser.*, 49:47–81, 1958.

[112] H. Rutishauser. Computational aspects of F. L. Bauer's simultaneous iteration method. *Numer. Math.*, 13:4–13, 1969.

[113] H. Rutishauser. Simultaneous iteration method for symmetric matrices. In J. H. Wilkinson and C. Reinsch, editors, *Handbook for Automatic Computation, Vol 2, Linear Algebra*, pages 284–302. Springer Verlag, Heidelberg – Berlin – New York, 1971.

[114] Y. Saad. On the rates of convergence of the Lanczos and the block Lanczos methods. *SIAM J. Numer. Anal.*, 17:687–706, 1980

[115] Y. Saad. Variations on Arnoldi's method for computing eigenelements of large unsymmetric matrices. *Linear Alg. Appl.*, 34:269–295, 1980.

[116] Y. Saad. *Numerical Methods for Large Eigenvalue Problems.* Manchester University Press, Manchester, UK, 1992.

[117] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Statist. Comput.*, 7:856–869, 1986.

[118] A. Sameh. On Jacobi and Jacobi-like algorithms for a parallel computer. *Math. Comp.*, 25:579–590, 1971.

[119] A. Schönhage. Zur Konvergenz des Jacobi-Verfahrens. *Numer. Math.*, 3:374–380, 1961.

[120] G. L. G. Sleijpen, J. G. L. Booten, D. R. Fokkema, and H. A. van der Vorst. Jacobi-Davidson type methods for generalized eigenproblems and polynomial eigenproblems. *BIT*, 36:595–633, 1996.

[121] G. L. G. Sleijpen and H. A. van der Vorst. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Applic.*, 17:401–425, 1996.

[122] G. L. G. Sleijpen, H. A. van der Vorst, and E. Meijerink. Efficient expansion of subspaces in the Jacobi-Davidson method for standard and generalized eigenproblems. *ETNA*, 7:75–89, 1998.

[123] G. L. G. Sleijpen, H. A. van der Vorst, and M. B. van Gijzen. Quadratic eigenproblems are no problem. *SIAM News*, 29:8–9, 1996.

[124] P. Smit. *Numerical Analysis of Eigenvalue Algorithms Based on Subspace Iterations.* PhD thesis, Tilburg University, The Netherlands, 1997.

[125] D. C. Sorensen. Implicit application of polynomial filters in a k-step Arnoldi method. *SIAM J. Matrix Anal. Applic.*, 13(1):357–385, 1992.

[126] G. W. Stewart. Simultaneous iteration for computing invariant subspaces of non-Hermitian matrices. *Numer. Math.*, 25:123–136, 1976.

[127] G. W. Stewart. SRRIT: A Fortran subroutine to calculate the dominant invariant subspace of a nonsymmetric matrix. Technical Report TR 154, Dept. of Comp. Science, University of Maryland, 1978.

[128] G. W. Stewart. On the early history of the singular value decomposition. *SIAM Rev.*, 35:551–566, 1993.

[129] G. W. Stewart and Ji-Guang Sun. *Matrix Perturbation Theory.* Academic Press, San Diego, CA, 1990.

[130] J. G. Sun. Perturbation analysis of quadratic eigenvalue problems. *BIT*, ?:??–??, 1999.

[131] O. Taussky. A recurring theorem on determinants. *Amer. math. Mon.*, 56:672–676, 1949.

[132] G. Temple. The accuracy of Rayleigh's method of calculating the natural frequencies of vibrating systems. *Proc. Roy. Soc. A*, 211:204–224, 1952.

[133] F. Tisseur. Backward error and condition of polynomial eigenvalue problems. Technical Report NA-332, University of Manchester, Manchester, UK, 1998.

[134] F. Tisseur and J. Dongarra. A parallel divide and conquer algorithm for the symmetric eigenvalue problem on distributed memory architectures. *SIAM J. Sci. Comput.*, 20:2223–2236, 1999.

[135] L. N. Trefethen. Pseudospectra of matrices. In D. F. Griffiths and G. A. Watson, editors, *Numerical Analysis 1991*, pages 234–266. Longman, 1992.

[136] L. N. Trefethen. Computation of pseudospectra. *Acta Numerica*, 8:247–295, 1999.

[137] L. N. Trefethen and D. Bau, III. *Numerical Linear Algebra.* SIAM, Philadelphia, PA, 1997.

[138] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. Hydrodynamic stability without eigenvalues. *Science*, 261:578-584, 1993.

[139] H. W. Turnbull and A. C. Aitken. *Introduction to the Theory of Canonical Matrices.* Dover Publications, NY, 1961.

[140] L. N. Trefethen and M. R. Trummer. An instability phenomenon in spectral methods. *SIAM J. Numer. Anal.*, 24:1008–1023, 1987.

[141] R. R. Underwood. An iterative block Lanczos method for the solution of large sparse symmetric eigenproblems. Technical Report STAN-CS-75-495, Dept. of Computer Science, Stanford University, Stanford, CA, 1975.

[142] A. van der Sluis and H. A. van der Vorst. The convergence behavior of Ritz values in the presence of close eigenvalues. *Linear Alg. Appl.*, 88/89:651–694, 1987.

[143] P. Van Dooren. The computation of Kronecker's canonical form of a singular pencil. *Linear Alg. Appl.*, 27:103–140, 1979.

[144] P. Van Dooren. The generalized eigenstructure problem in linear system theory. *IEEE Trans. Auto. Cont. AC*, 26:111–128, 1981.

[145] P. Van Dooren. A generalized eigenvalue approach of solving Ricatti equations. *SIAM J. Sci. Statist. Comput.*, 2:121–135, 1981.

[146] P. Van Dooren. Algorithm 590,DUSBSP AND EXCHQZZ: FORTRAN subroutines for computing deflating subspaces with specified spectrum. *ACM Trans. Math. Softw.*, 8:376–382, 1982.

[147] H. P. M. van Kempen. On quadratic convergence of the special cyclic Jacobi method. *Numer. Math.*, 9:19–22, 1966.

[148] R. S. Varga. *Matrix Iterative Analysis.* Prentice-Hall, Englewood Cliffs N.J., 1962.

[149] R. von Mises and Hilda Pollaczek-Geiringer. Praktische Verfahren der Gleichungsauflösung. Zeitschrift für Angewandte Mathematik und Mechanik, 9:58–77,152–164, 1929.

[150] D. S. Watkins. *Fundamentals of Matrix Computations.* John Wiley and Sons, New York, 1991.

[151] D. S. Watkins. Some perspectives on the eigenvalue problem. *SIAM Rev.*, 35:430–471, 1993.

[152] D. S. Watkins and L. Elsner. Chasing algorithms for the eigenvalue problem. *SIAM J. Matrix Anal. Applic.*, 12:374–384, 1991.

[153] H. F. Weinberger. *Variational Methods for Eigenvalue Approximation.* SIAM, Philadelphia, 1974.

[154] J. H. Wilkinson. *The Algebraic Eigenvalue Problem.* Clarendon Press, Oxford, 1965.

[155] J. H. Wilkinson. Linear differential equations and Kronecker's canonical form. In C. de Boor and G. H. Golub, editors, *Recent advances in Numerical Analysis*, pages 231–265, New York, 1978. Academic Press.

[156] J. H. Wilkinson. Kronecker's canonical form and the QZ algorithm. *Linear Alg. Appl.*, 28:285–303, 1979.

[157] J. H. Wilkinson. Eigenvalue problems. In A. Iserles and M. J. D. Powell, editors, *The State of the Art in Numerical Analysis*, pages 1–39, Oxford, 1987. Clarendon Press.

[158] J. H. Wilkinson and C. Reinsch, editors. *Handbook for Automatic Computation, Vol 2, Linear Algebra.* Springer Verlag, Heidelberg – Berlin – New York, 1971.

[159] T. Zhang, K. H. Law, and G. H. Golub. On the homotopy method for perturbed symmetric generalized eigenvalue problems. *SIAM J. Numer. Anal.*, 35:1, 300–319, 1998.