

STICHTING
MATHEMATISCH CENTRUM
2e BOERHAAVESTRAAT 49
AMSTERDAM

MR 107 *12*

Calculation of eigenvalues and eigenvectors.

(Informatie, 11(1969), p 118-124.)

by

T.J. Dekker



1969

CALCULATION OF EIGENVALUES AND EIGENVECTORS

by dr. T. J. Dekker

ABSTRACT

This paper is a written version of a talk delivered by the author before the Nederlands Rekenmachine Genootschap (Dutch Association for Computing Machinery) on May 28, 1968, at Amsterdam.

A survey is given of the most popular and successful numerical methods for calculating the eigenvalues and eigenvectors of full real matrices. Attention is focused on methods for general matrices; methods for symmetric matrices are treated only as special cases. The successive topics are: an introduction with an outline of the theory; reducing similarity transformations, in particular, Wilkinson's and Householder's transformation to Hessenberg form; methods to calculate eigenvalues of Hessenberg matrices, viz. the QR method and nondeflating methods using Hyman's formula; methods to calculate eigenvectors of Hessenberg matrices, viz. inverse iteration and the QR method; Osborne's equilibration; Eberlein's generalized Jacobi process.

1 INTRODUCTION

During the last decade, the art of computing eigenvalues and eigenvectors of square matrices has made considerable progress. Numerical methods nowadays are available which work satisfactorily for most matrices occurring in practice. For real symmetric (and complex Hermitian) matrices in particular, we have quite stable numerical processes whose convergence can be guaranteed.

We shall try to give a survey of the most popular and successful methods for calculating eigenvalues and eigenvectors, restricting ourselves to methods for full (i.e., nonsparse) real square matrices. We shall pay attention mainly to the general problem in which the matrix is not assumed to be symmetric, and deal with methods for symmetric matrices only as special cases of the general methods.

To define the problem and some related notions, let M be a matrix of the order n . The „eigenvalues” of M are those real or complex numbers λ for which the linear system $Mx = \lambda x$ has a nonnull solution vector x , called „eigenvector” of M corresponding to λ . The eigenvalues of M are the zeros of $\det(M - \lambda I)$, which is a polynomial of degree n in λ , the „characteristic polynomial” of M . Hence, M has at most n distinct eigenvalues, and at least one eigenvector corresponds to each eigenvalue. A „similarity transformation” is a transformation which, to each n th order matrix M , associates the matrix $S^{-1}MS$, where S is any nonsingular n th order matrix. The matrices M and $S^{-1}MS$ are called „similar”. Similarity transformations are important, because they leave the eigenvalues invariant, and transform the eigenvectors in the following simple way: if x is an eigenvector of $S^{-1}MS$ corresponding to the eigenvalue λ , then Sx is an eigenvector of M corresponding to λ .

Matrix M is called „diagonalizable”, if it is similar to a diagonal matrix: $M = X\Lambda X^{-1}$. The eigenvalues of M are the diagonal elements of the similar diagonal matrix

Λ , the columns of the transforming matrix X are the corresponding eigenvectors of M , and the rows of X^{-1} are the corresponding eigenvectors of the transposed matrix M^T . If M is diagonalizable and has multiple eigenvalues, then the eigenvectors of M and M^T are not unique, but may always be chosen (and are preferably chosen) such that the matrix of eigenvectors of M and M^T are each other's inverse. (The eigenvectors of M and M^T then form a „biorthogonal” system.) If M has only simple eigenvalues, then it is diagonalizable, and its eigenvectors are unique up to a scalar factor.

Computing eigenvalues and eigenvectors is considerably simpler for real symmetric (and complex Hermitian) matrices than for other ones, because a real symmetric (complex Hermitian matrix) has the following nice properties:

- 1 all eigenvalues are real;
- 2 the matrix is diagonalizable;
- 3 the eigenvectors can be chosen such that the matrix of eigenvectors is real orthogonal; (3': complex unitary).

Nearly as easy to handle are the normal matrices, which are characterized by properties (2) and (3'); in other words, a matrix is normal if it can be written in the form $U\Lambda U^{-1}$, where Λ is diagonal and U unitary.

The eigenvalues of normal matrices are well conditioned; i.e., small changes in the elements of a normal matrix cause small changes in the eigenvalues. Eigenvectors are ill conditioned if they correspond to closely clustered eigenvalues. If two eigenvalues coincide, then any linear combination of two (linearly) independent corresponding eigenvectors is again an eigenvector; this explains why two eigenvectors corresponding to close eigenvalues are very sensitive to perturbations of the elements of the matrix.

The computation of eigenvalues and eigenvectors of nonnormal matrices may be complicated for the following two reasons.

- 1 If the matrix has multiple eigenvalues, then it may be nondiagonalizable. A nondiagonalizable n th order matrix does not have n linearly independent eigenvectors (its eigensystem is „defective”). The notion „eigenvector” may be generalized to „principal vector” corresponding to an eigenvalue λ , i.e., a nonnull vector, x , satisfying

$$(M - \lambda I)^j x = 0,$$

where the integer j , the „grade” of x , is chosen as small as possible. If M is diagonalizable, then all its principal vectors are of grade 1, i.e., are ordinary eigenvectors. The computation of principal vectors of grade > 1 is a nearly unexplored area (see, however, J. M. Varah [14]), and will not be discussed here.

- 2 The eigenvalues of a matrix may be ill conditioned, even if the matrix is diagonalizable. Let $M = X\Lambda X^{-1}$, where Λ is diagonal. Let $\|X\|$ denote the „spectral norm” of X , i.e., the square root of the

largest eigenvalue of $\bar{x}^T x$. Then we have the following:

Theorem (Bauer and Fike [4] [1, p. 87])

If λ is any eigenvalue of the perturbed matrix $M + \delta M$, then there is an eigenvalue λ_i of M such that

$$|\lambda_i - \lambda| < \|X^{-1}\| \|X\| \|\delta M\|.$$

In other words, the minimum, μ , of $\|X^{-1}\| \|X\|$ for all permissible eigenvector matrices X of M is a measure for the sensitivity of the eigenvalues to perturbations of the elements of M . If M is normal, then $\mu = 1$ and the eigenvalues are well conditioned. On the other hand, if M is far from normal (especially if M is rather close to a nondiagonalizable matrix), then μ is considerably larger than 1, and the eigenvalues (or at least some of them) are ill conditioned. More specifically, a measure for the sensitivity of a (simple) eigenvalue λ of M is the quantity

$$\frac{\|y\| \|x\|}{|y^T x|},$$

where x is a corresponding eigenvector of M and y of M^T , and where $\|x\|$ denotes the Euclidean vector norm.

Besides the spectral norm mentioned above, one uses the following matrix norms, which are much easier to calculate:

the Euclidean matrix norm

$$\|M\|_E = \sqrt{\sum_{i,j=1}^n |M_{ij}|^2}$$

and the infinity norm

$$\|M\|_\infty = \max_i \sum_{j=1}^n |M_{ij}|.$$

In particular, we use matrix norms to define realistic tolerances for the various stages of calculating eigenvalues and eigenvectors. These tolerances have the form: matrix norm times some specified dimensionless parameter (e.g., the machine precision).

The theory of matrices and methods to compute eigenvalues and eigenvectors are extensively treated in J. H. Wilkinson [1], A. S. Householder [2] and D. K. Faddeev and V. N. Faddeeva [3]. In particular, [1] deals with various numerical methods from a practical standpoint, including error analyses and assessments of the methods based on the practical experience of the author. From this book, we have taken a great deal of the material presented here. Our survey is certainly not complete. The research on this subject is very active, and there are several interesting developments which we shall not discuss, in particular, the calculation of bounds for a computed eigensystem [14], and methods to improve approximate eigensystems [1, p. 637-646].

2 REDUCING SIMILARITY TRANSFORMATIONS

Many methods for calculating the eigenvalues and/or eigenvectors of a matrix start with a similarity transformation reducing the matrix to one having a special form which is much easier to handle. Some important special forms obtainable by means of a direct, i.e., noniterative, process are the Hessenberg form, the tridiagonal form and the Frobenius form.

A matrix, H , is „upper-Hessenberg” or „almost upper-triangular” if the elements below its first subdiagonal are zero; i.e., $H_{ij} = 0$ for $i > j + 1$. A matrix, T , is „tridiagonal,” if all elements outside the main diagonal and the adjacent codiagonals are zero; i.e., $T_{ij} = 0$ for $|i - j| > 1$.

A matrix is „Frobenius,” if it is the direct sum of elementary Frobenius matrices, which are of the form

$$F = \begin{bmatrix} b_{r-1} & & & & b_0 \\ .1 & & 0 & & \\ & .0 & & & \\ \vdots & & & & \\ & .0 & & & 0 & 1 & 0 \end{bmatrix}$$

The characteristic polynomial of F is

$$\lambda^r - b_{r-1}\lambda^{r-1} \dots - b_1\lambda - b_0;$$

therefore F is also called the „companion matrix” of this polynomial.

2.1 Reduction to Hessenberg Form

For reducing a general matrix to a similar upper-Hessenberg matrix, there exist completely satisfactory, stable transformations. We discuss two of them which are most commonly used, viz. Wilkinson's and Householder's transformation.

Wilkinson's transformation [1, p. 357-368] is a triangular transformation with stabilizing interchanges. The transforming matrix is the product of a permutation matrix, P , and a unit lower-triangular matrix, L („unit” means that the diagonal elements are 1). Thus, if M is the given matrix and H the resulting upper-Hessenberg matrix, then we have

$$H = L^{-1}P^{-1}MPL.$$

The interchanges are chosen such that $|L_{ij}| < 1$ for all i and j . Moreover, one can choose the elements L_{ij} , $i > 2$ arbitrarily; in practice, one chooses them equal to 0. The process resembles the triangular decomposition used to solve linear systems. The transformation is more accurate if the scalar products involved are calculated in extra precision. The number of operations (each operation consisting of a multiplication and an addition) roughly equals $\frac{1}{2}n^3$ for large n .

In exceptional cases, the transformation might lead to growing of the pivots (i.e., the elements chosen in each step which are going to be the subdiagonal elements of the resulting Hessenberg matrix) by a factor up to 2^{n-2} (in a similar way as might occur in triangular decomposition with partial pivoting [1, p. 212]). Furthermore, matrix L might be ill conditioned. Either would cause a substantial loss of accuracy. No analysis of these phenomena appears to have been made; in practice, Wilkinson's transformation turns out to be very satisfactory.

Householder's transformation [2, p. 162] [5] [1, p. 290-299, p. 347-353] is an orthogonal transformation. The transforming matrix, P , is the product of $n - 2$ Householder matrices, which are orthogonal symmetric matrices of the form

$$I + kuu^T,$$

where u is a column vector and k a scalar satisfying $ku^T u = -2$ (this condition insures orthogonality of the matrix). In the j th Householder matrix, $j = 1, \dots, n-2$, the first j elements of u are zero, and the other elements of u are chosen such that the desired zeros are introduced in the j th column of the matrix. Although part of the calculations involves scalar products, only marginally more accuracy is obtained if the scalar products are calculated in extra precision. The number of operations (i.e., multiplications and additions) roughly equals $\frac{5}{3} n^3$ for large n , so Householder's transformation is a factor two slower than Wilkinson's. On the other hand, Householder's transformation is completely stable. For a symmetric matrix Householder's transformation has the great advantage that symmetry is preserved (because the transformation is orthogonal); so the amount of work is considerably reduced, the number of operations roughly being equal to $\frac{3}{2} n^3$ for large n . The resulting similar matrix is tridiagonal, since it is symmetric and Hessenberg.

2.2 Further Reductions

After reducing the given matrix to a similar upper-Hessenberg matrix, one may perform a further reduction to tridiagonal or Frobenius form.

A tridiagonal matrix similar to an upper-Hessenberg matrix may be obtained by means of the transposed Wilkinson process without interchanges; i.e., the transforming matrix is a unit upper-triangular matrix, U , and the resulting matrix has lower-Hessenberg form. Moreover, the upper-Hessenberg form turns out to be preserved, so that the resulting similar matrix is tridiagonal. Alas, stabilizing interchanges cannot be used, because they would destroy the upper-Hessenberg form. Using no interchanges, however, the process may be unstable, and, in fact, breaks down if a zero pivot occurs. One may try to avoid this by choosing another first column of U . No satisfactory process seems to be known. In practice, it is often advisable to use double (or multiple) precision in the reduction to tridiagonal form. We shall not go further into this [1, p. 395-404].

The reduction of an Hessenberg matrix to Frobenius form amounts to calculating the coefficients of the characteristic polynomial. This method is often very unstable as the eigenvalues may be much more sensitive to errors in the coefficients of the characteristic polynomial than to errors in the elements of the given Hessenberg matrix. In general, the reduction to Frobenius form is much less satisfactory than the reduction to tridiagonal form [1, p. 405-411].

3 EIGENVALUES OF HESSENBERG MATRICES

There exist two main types of methods for calculating eigenvalues, viz. deflating and nondeflating methods. In a deflating method, the order of the matrix involved is decreased as soon as an eigenvalue has been found; in a nondeflating method, the order of the matrix involved remains constant.

3.1 A Deflating Method: QR Iteration

The most satisfactory method, nowadays, for finding the eigenvalues of an upper-Hessenberg matrix is J. G. F. Francis' QR iteration [6] [1, p. 515 sqq] [7] [8] [14]. Let $H^{(0)}$ be a given (not necessarily upper-Hessenberg)

matrix. For $k = 0, 1, 2, \dots$, we first calculate the decomposition

$$(3.1.1) \quad H^{(k)} = Q^{(k)} R^{(k)},$$

where $Q^{(k)}$ is orthogonal and $R^{(k)}$ upper-triangular. This decomposition may be obtained mathematically (the numerical computation uses a different formula) by means of Gram-Schmidt orthogonalization, which is performed in n stages as follows. (We drop the superscripts „(k)“.)

In the j th stage, $j = 1, \dots, n$, the j th columns of Q and R are calculated according to the formulas

$$(3.1.2) \quad \left\{ \begin{array}{l} R_{ij} = \sum_{k=1}^{j-1} Q_{ki} H_{kj}, \quad i = 1, \dots, j-1, \\ \tilde{Q}_{ij} \text{ (say)} = H_{ij} - \sum_{k=1}^{j-1} Q_{ik} R_{kj}, \quad i = 1, \dots, n, \\ R_{jj} = \left(\sum_{i=1}^n \tilde{Q}_{ij}^2 \right)^{1/2} \\ Q_{ij} = \tilde{Q}_{ij} / R_{jj}, \quad i = 1, \dots, n. \end{array} \right.$$

From these formulas it easily follows that the decomposition (3.1.1) is always possible, and, if H is nonsingular and one imposes the extra condition that the diagonal elements of R be positive, then also unique; moreover, it is obvious that, if H is upper-Hessenberg, then also Q .

After calculating the decomposition (3.1.1), we subsequently form the product

$$(3.1.3) \quad H^{(k+1)} = R^{(k)} Q^{(k)}.$$

Thus, $H^{(k+1)}$ is obtained from $H^{(k)}$ by an orthogonal similarity transformation:

$$(3.1.4) \quad H^{(k+1)} = (Q^{(k)})^{-1} H^{(k)} Q^{(k)}.$$

Moreover, if $H^{(k)}$ is upper-Hessenberg, then also $H^{(k+1)}$, since $R^{(k)}$ is upper-triangular and $Q^{(k)}$ upper-Hessenberg.

The sequence thus obtained satisfies the following:

Theorem (Francis [6])

If $H^{(0)}$ is a nonsingular matrix having eigenvalues of distinct modulus, then the sequence of matrices $H^{(k)}$ converges to an upper-triangular matrix.

From now on, we assume that $H^{(0)}$ is upper-Hessenberg. If $H^{(0)}$ can be diagonalized and its subdiagonal elements

$H_{i+1,i}^{(0)}$, $i = 1, \dots, n-1$, are nonzero, then $H^{(0)}$ has no multiple eigenvalues. Moreover, we have:

Theorem (Francis [6])

If $H^{(0)}$ is a nonsingular diagonalizable upper-Hessenberg matrix whose subdiagonal elements are nonzero, and whose eigenvalues are of distinct modulus, then the main diagonal elements of $H^{(k)}$ converge to the eigenvalues in order of magnitude.

Shift Strategy

The elements $H_{nn-1}^{(k)}$ converge linearly to 0, the con-

vergence factor being $|\lambda_n|/|\lambda_{n-1}|$, where λ_n is the eigenvalue of smallest modulus and λ_{n-1} the next smallest. So, this convergence is usually too slow. In order to speed up convergence, one uses a „shift” strategy, i.e., the above formulas (3.1.1) and (3.1.3) are modified as follows:

$$(3.1.5) \quad \begin{cases} H^{(k)} - s^{(k)}I = Q^{(k)}R^{(k)}, \\ H^{(k+1)} = R^{(k)}Q^{(k)} + s^{(k)}I, \end{cases}$$

where $s^{(k)}$ is a suitably chosen shift. In practice, one chooses either

$$(3.1.6) \quad s^{(k)} = H_{nn}^{(k)}$$

or preferably

$$(3.1.7) \quad s^{(k)} = \mu^{(k)},$$

i.e., the eigenvalue of the lower right hand two-by-two submatrix of $H^{(k)}$ closest to $H_{nn}^{(k)}$. (Francis proposes to

perform the first iteration steps with shift $s^{(k)} = 0$, and to start using the shift $s^{(k)} = \mu^{(k)}$, when $\mu^{(k)}$ has become invariant within one binary digit precision; this strategy makes it more likely that the eigenvalues will be found in order of increasing modulus.)

These shift strategies yield a process which, if it is convergent converges nearly always quadratically for diagonalizable matrices (or cubically for symmetric matrices). The choice $s^{(k)} = \mu^{(k)}$ mentioned above may turn out to be nonreal. In that case the shift strategy may be modified in one of the three following ways:

- 1 for real matrices having only real eigenvalues (in particular real symmetric matrices), one may choose $s^{(k)} = \text{real part of } \mu^{(k)}$;
- 2 for complex (or real) matrices, one may choose $s^{(k)} = \mu^{(k)}$, and use complex arithmetic whenever necessary [6] [7].

This „single” QR iteration process requires roughly $12n^2$ real multiplications per step.

- 3 for real matrices having complex eigenvalues the previous strategy is possible, but expensive. A much faster process is obtained by choosing $s^{(k)} = \mu^{(k)}$, $s^{(k+1)} = \bar{\mu}^{(k)}$.

This „double” QR iteration process [6] [8] [14] has the advantage that, from a real iterate $H^{(k)}$, one obtains a real iterate $H^{(k+2)}$ which can be calculated in real arithmetic requiring only $5n^2$ real multiplications per double step.

Deflation

As soon as, for some i , $|H_{ii-1}^{(k)}|$ has become small enough (e.g., $<$ some norm of $H^{(0)}$ times the machine precision), this element can be neglected so that $H^{(k)}$ approximately equals

$$(3.1.8) \quad \begin{bmatrix} H_1 & M \\ 0 & H_2 \end{bmatrix}$$

where H_1 is of the order i , H_2 of the order $n-i$, 0 is the $n-i$ by i null matrix and M an i by $n-i$ matrix. The eigenvalues of this matrix are equal to those of H_1 and H_2 . Thus, the problem of finding the eigenvalues of $H^{(0)}$ is reduced to two problems involving smaller matrices.

In particular, if $H_{nn-1}^{(k)}$ or $H_{n-1, n-2}^{(k)}$ is negligible, then H_2 is of order 1 or 2, and its eigenvalues can be calculated directly; thus, the problem is reduced to finding the eigenvalues of H_1 of order $n-1$ or $n-2$. The process is completed if the successive reductions have produced matrices of order 1 or 2 only.

Convergence

Since a matrix H_2 of order 1 or 2 can be handled directly, one calls a QR iteration process „convergent” if either

$H_{nn-1}^{(k)}$ or $H_{n-1, n-2}^{(k)}$ converges to 0. Using the shifts described above, especially (3.1.7) and the modifications, the process converges in the great majority of cases, but, unfortunately, there exist matrices for which no convergence occurs; e.g. the n th order permutation matrix ($n > 2$)

$$(3.1.9) \quad \begin{bmatrix} 0 & & & & 0 & 1 \\ & \dots & & & & 0 \\ 1 & & & & & \\ 0 & & & & & \\ \vdots & & & & & \\ 0 & & 0 & 1 & \dots & 0 \end{bmatrix}$$

is invariant under the QR transformation with zero shift; moreover, according to either shift strategy mentioned above, the shift equals 0; so obviously no convergence occurs. One might choose a shift $s^{(k)} = \mu^{(k)} + \alpha H_{nn-1}^{(k)}$ where α is suitably chosen [1, p. 511]

[8]. Using this shift strategy, the process is convergent for matrix (3.1.9) and, probably for the great majority of matrices, but presumably for any constant α there are other matrices for which this shift strategy will fail.

For symmetric tridiagonal matrices, the iterates $H^{(k)}$ remain symmetric tridiagonal, since each step is an orthogonal similarity transformation (3.1.4); moreover, the shift strategies (3.1.6) and (3.1.7) both yield a convergent process [9].

Advantages

- 1 As the QR step is an orthogonal similarity transformation (3.1.4), the condition number μ (see page 119) of the eigenvalues of the matrix remains invariant. This property makes the process numerically very stable.
- 2 Deflation reduces the problem to problems involving matrices of lower order. Moreover, after deflation, the matrix H_1 involved is usually closer to the limit than an arbitrary upper-Hessenberg matrix of the same order.
- 3 Nonconvergence seldom occurs. According to J. H. Wilkinson [1, p. 538] the double QR iteration, with shifts according to (3.1.7) and deflation, is the „most powerful general-purpose programme of any I have used.”

3.2 Nondeflating Methods

The characteristic polynomial of an upper-Hessenberg matrix H may be evaluated according to Hyman's formula as follows [10] [1, p. 426-429]. We calculate a vector x such that all components of $(H-\lambda I)x$ except the first are zero and $x_n = 1$. We then find

$$(3.2.1) \quad x_{i-1} = \left((H_{ii}-\lambda)x_i + \sum_{j=i+1}^n H_{ij}x_j \right) / H_{i,i-1},$$

$$i = n, n-1, \dots, 2,$$

and obtain the value of the characteristic polynomial of H at λ according to

$$(3.2.2) \quad \det(H-\lambda I) = \left((H_{11}-\lambda)x_1 + \sum_{j=2}^n H_{1j}x_j \right) \prod_{j=2}^n (-H_{jj-1}).$$

If some of the subdiagonal elements $H_{i,i-1}$ are 0 or negligible, we may either partition the matrix as above (3.1.8) and handle H_1 and H_2 separately, or replace these negligible elements by some threshold (e.g., some norm of H times the machine precision).

Thus, we may assume that the subdiagonal elements are nonzero, and, since we are searching for zeros of the characteristic polynomial, we may disregard the nonzero factor

$$\prod_{j=2}^n (-H_{jj-1}).$$

Therefore we replace (3.2.2) by

$$(3.2.3) \quad f(\lambda) = (H_{11}-\lambda)x_1 + \sum_{j=2}^n H_{1j}x_j,$$

and search for the zeros of $f(\lambda)$, or rather, after finding the eigenvalues $\lambda_1, \dots, \lambda_{i-1}$, we search for a zero of

$$(3.2.4) \quad f_i(\lambda) = f(\lambda) / \prod_{j=1}^{i-1} (\lambda - \lambda_j).$$

The zeros of $f(\lambda)$ or $f_i(\lambda)$ may be calculated by means of any standard iteration process [1, p. 435-461]. If all eigenvalues are real, then one may use linear interpolation, Newton or Laguerre. To find complex eigenvalues, one may use quadratic interpolation (formula of Muller [1, p. 435] or Traub [1, p. 484]), Bairstow (adapted to polynomials in the form (3.2.3) [1, p. 449-451]) and Laguerre [11] [12].

The derivatives needed for Newton and Laguerre iteration are obtained by differentiating (3.2.1) and (3.2.3); the resulting formulas are rather similar, and the number of operations for calculating $f'(\lambda)$ or $f''(\lambda)$ is nearly the same as for calculating $f(\lambda)$.

Advantages and Disadvantages

- 1 In the whole process, the original matrix H is used, so that we may expect to obtain more accurate eigenvalues than in the QR method, especially if the latter requires many iterations.
- 2 Since no deflation is performed, clusters of eigenvalues cause trouble: numerator and denominator in (3.2.4) both may become 0 or very small, so that

no figure of $f_i(\lambda)$ is found correctly; convergence is merely linear and often slow; the multiplicity of the eigenvalues may be difficult to determine; sometimes it is not known whether a calculated eigenvalue is real or not, so that one cannot properly divide the eigenvalues into pairs of complex conjugate eigenvalues and real ones.

- 3 No iteration process is known which converges for all (Hessenberg) matrices. The most successful seems to be Laguerre's formula.
- 4 Since no deflation is performed, the process is considerably slower than QR iterations; e.g., iteration according to Laguerre's formula turns out to be about three times slower (on the Electrologica X8) than the QR method.

Symmetric Matrices

Symmetric tridiagonal matrices whose codiagonal elements are nonzero possess the important property that the determinants of the principal minors (the i th principal minor being the submatrix consisting of the first i rows and columns) form a Sturm sequence; the signs of these determinants completely determine the number of eigenvalues smaller (or larger) than the argument at which these determinants were evaluated.

Using this property one can calculate the eigenvalues by means of bisection, preferably combined with linear interpolation (or one of the other processes mentioned above) [1, p. 299-315] [8]. These methods for symmetric matrices certainly do not have the disadvantages mentioned above for the general case, but, in fact, are competitive with QR iteration.

4 EIGENVECTORS OF HESSENBERG MATRICES

Eigenvectors may be calculated by means of inverse iteration or, if the QR method is used for calculating the eigenvalues, by means of a direct method using the results of the QR iteration. After calculating the eigenvectors of a Hessenberg matrix H , one obtains the eigenvalues of the similar matrix $M = SHS^{-1}$ by means of back transformation: if x is an eigenvector of H , then Sx is the corresponding eigenvector of M .

4.1 Inverse Iteration

Let H be a given (not necessarily Hessenberg) matrix of the order n , λ an approximate eigenvalue of H , and $x^{(0)}$ an n -vector. For $k = 0, 1, 2, \dots$, we calculate $y^{(k)}$ by solving the linear system

$$(4.1.1) \quad (H-\lambda I)y^{(k)} = x^{(k)}$$

(by means of Gaussian elimination with row interchanges), and obtain $x^{(k+1)}$ by normalizing $y^{(k)}$:

$$(4.1.2) \quad x^{(k+1)} = \frac{1}{\|y^{(k)}\|} y^{(k)}.$$

If H is diagonalizable and x_1, \dots, x_n are n linearly independent eigenvectors of H , then we can write

$$x^{(k)} = \alpha_1 x_1 + \dots + \alpha_n x_n.$$

Hence

$$y^{(k)} = \frac{\alpha_1}{\lambda_1 - \lambda} x_1 + \dots + \frac{\alpha_n}{\lambda_n - \lambda} x_n.$$

So, if λ_1 is well separated from the other eigenvalues, λ is a reasonable approximation to λ_1 , and α_1 does not (nearly) vanish, then the sequence of vectors $x^{(k)}$ converges (rapidly) to the eigenvector x_1 . In practice, one or two iteration steps nearly always suffice. If H is a diagonalizable matrix having multiple (or closely clustered) eigenvalues, and one wants to find linearly independent corresponding eigenvectors, then one, obviously, cannot use the same numerical approximation λ twice in the inverse iteration. Wilkinson remarked [1, p. 328] that, in this case, the inverse iteration is very sensitive to small changes in λ , so that replacing λ by another nearby value and performing another inverse iteration, one finds an approximate eigenvector which is not almost linearly dependent from the previously calculated one(s).

If H is symmetric, then one obtains a set of orthogonal eigenvectors by calculating the component orthogonal to the eigenvectors already found (here one needs to consider only those eigenvectors which belong to close eigenvalues). This orthogonal component is calculated in each inverse iteration step and used as starting vector $x^{(k)}$ for the next step [8]. If H is not symmetric, one can similarly obtain a biorthogonal set of eigenvectors of H and H^T .

If H is real and λ complex nonreal, then we have the following four possible strategies [1, p. 629-630].

- 1 We may use complex arithmetic where necessary. This strategy is implemented in [8].
- 2 We may work entirely with real arithmetic as follows.
We write $\lambda = \xi + i\eta$, $x^{(k)} = p^{(k)} + iq^{(k)}$
 $y^{(k)} = u^{(k)} + iv^{(k)}$, and equate real and imaginary parts in (4.1.1). This gives

$$(4.1.3) \quad (H - \xi I)u^{(k)} + \eta v^{(k)} = p^{(k)},$$

$$(4.1.4) \quad -\eta u^{(k)} + (H - \xi I)v^{(k)} = q^{(k)}.$$

After solving this real system of the order $2n$, we normalize $u^{(k)} + iv^{(k)}$ and obtain $p^{(k+1)} + iq^{(k+1)}$. This strategy requires about twice as much space and time, and is therefore not acceptable.

- 3 From (4.1.3) and (4.1.4) we can easily derive the equations

$$(4.1.5) \quad ((H - \xi I)^2 + \eta^2 I)u^{(k)} = (H - \xi I)p^{(k)} - \eta q^{(k)},$$

$$(4.1.6) \quad ((H - \xi I)^2 + \eta^2 I)v^{(k)} = \eta p^{(k)} + (H - \xi I)q^{(k)}.$$

Wilkinson remarks [1, p. 630] that using these formulas in the inverse iteration step, one does not obtain a good approximation to an eigenvector.

- 4 One may calculate $u^{(k)}$ by solving (4.1.5), and $v^{(k)}$ from (4.1.3), and then normalize to produce $p^{(k+1)}$ and $q^{(k+1)}$. This strategy requires more space (viz. for storing the matrix H^2) but slightly less time than strategy (1), and is implemented by Varah in [13] [14].

4.2 Calculating Eigenvectors Using QR Iteration

The QR iteration produces an upper-triangular matrix U which, apart from rounding errors, is similar to the given matrix H . Thus, we have

$$H = QUQ^{-1}$$

where Q is the product of the rotations $Q^{(k)}$ and the rotations transforming the two-by-two blocks into pairs of one-by-one blocks. Let H , and thus also U , be diagonalizable; let T be a matrix of eigenvectors of U . Then, obviously, QT is a matrix of eigenvectors of H . The eigenvectors of U may be chosen such that T is unit upper-triangular, and are calculated by solving the corresponding triangular linear systems. Thus, we obtain

$$T_{ij} = \left(\sum_{k=i+1}^j U_{ik} T_{kj} \right) / (U_{ji} - U_{ii}), \quad i < j.$$

If the denominator (nearly) vanishes (i.e., if the i th and j th eigenvalues are (nearly) equal), then it should be replaced by a tolerance (e.g., some norm of H times the machine precision), and, since we assume that U is diagonalizable, then also the numerator (nearly) vanishes, so that we obtain eigenvectors which are not almost linearly dependent.

This process is implemented in [8] and turns out to be competitive with inverse iteration as to precision obtained and time required; for some matrices, the eigenvectors obtained were more distant from a linearly dependent set than the eigenvectors produced by inverse iteration. For a symmetric matrix H , the matrix U produced by the QR iteration is, of course, diagonal. Hence, the matrix X of eigenvectors of H simply equals Q . If S is the transformation matrix transforming the given M into H (i.e., $H = S^{-1}MS$), then SQ is the matrix of eigenvectors of M . Starting from S , the matrix SQ can be built up during the QR iteration by multiplying S by the matrix $Q^{(k)}$ in each step. Thus no back transformation is needed afterwards and memory space is saved [15] [8].

5 OSBORNE'S EQUILIBRATION

A matrix is called „equilibrated” with respect to its eigenproblem, if each row has the same Euclidean norm as the corresponding column. E. E. Osborne [16] proved that, if M is irreducible, i.e., if there exists no permutation matrix P such that

$$P^{-1}MP = \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix}$$

where 0 is a null matrix, then there exists a nonsingular diagonal matrix D , such that $D^{-1}MD$ is equilibrated. An equilibrated matrix has the property that, in the class of matrices which are similar to it by a diagonal similarity transformation, it has minimum Euclidean norm. It turns out that, for calculating eigenvalues and eigenvectors, it is of considerable importance to work with a (roughly) equilibrated matrix, for the following two reasons.

- 1 To determine useful tolerances for the various stages of the calculation, it is desirable to have a not too severe overestimate of the spectral radius (i.e., the maximum absolute value of the eigenvalues) of the matrix. The Euclidean norm of an equilibrated matrix is an acceptable estimate because of the minimum property mentioned above; on the other hand, if the matrix is far from equilibrated („badly balanced”), then the norm is a severe overestimate.

- 2 If orthogonal transformations are used (Householder's transformation and/or QR iteration), then badly balanced matrices should be avoided, because the calculations (especially the QR iteration) are much stabler for (roughly) equilibrated matrices. On the other hand, Wilkinson's transformation and inverse iteration (apart from the choice of the interchanges involved), and also Hyman's formula, are „scaling invariant” (i.e., invariant under diagonal similarity transformations).

Osborne showed that equilibration, i.e., transformation of a matrix into an equilibrated matrix by means of a diagonal similarity transformation, can be performed by means of an iterative process; each step consists of a diagonal similarity transformation equating the Euclidean norm of a certain row and corresponding column. If the matrix is irreducible, then the process converges and the convergence is usually very fast. In practice, one chooses the diagonal elements of the transforming matrices equal to a power of two, in order to maintain exact similarity.

If the given matrix is symmetric, equilibration is not needed, because symmetric matrices are equilibrated by definition.

6 EBERLEIN'S METHOD

Among the methods which do not start with a reducing similarity transformation, Eberlein's method [17] [18] [19] seems to be very promising. This method is a generalization of Jacobi's method; by means of successive similarity transformations, the given matrix, M , is transformed into a matrix whose Euclidean norm is minimal. The method may start with Osborne's equilibration, which uses only diagonal transformations to minimize the Euclidean norm. In subsequent steps of Eberlein's process, plane transformations (i.e., transformations affecting only two rows and columns of the matrix) are performed.

In each step, the Euclidean norm decreases, and hence also the „departure from normality” (i.e., the squared Euclidean norm minus the sum of the squared moduli of the eigenvalues), which vanishes for normal matrices. The process first converges to a normal matrix; subsequently, plane (complex) rotations (i.e., plane unitary transformations) are performed. The process converges to a block diagonal matrix whose blocks are nearly always one-by-one or two-by-two, so that the eigenvalues and eigenvectors are then obtained immediately. For symmetric matrices (and essentially also for normal matrices), Eberlein's process reduces to Jacobi's process, in which only plane (complex) rotations are performed. Advantages of Eberlein's method are:

- 1 Convergence of the (mathematical) process can be proved.
- 2 The accuracy of the results is acceptable (comparable to QR).
- 3 It is more economical in storage than QR with inverse iteration, and as economical as QR with direct calculation of the eigenvectors.

A disadvantage is that Eberlein's method is slower than QR.

REFERENCES

- 1 Wilkinson, J. H., The algebraic eigenvalue problem, Clarendon Press, Oxford 1965.
- 2 Householder, A. S., The theory of matrices in numerical analysis, Blaisdell Publishing Co., New York 1964.
- 3 Faddeev D. K. and V. N. Faddeeva, Computational methods of linear algebra (translated from Russian), Freeman and Co., San Francisco 1963.
- 4 Bauer F. L. and C. T. Fike, Norms and exclusion theorems, Num. Mat. 2 (1960), 137-141.
- 5 Householder A. S. and F. L. Bauer, On certain methods for expanding the characteristic polynomial, Num. Mat. 1 (1959), 29-37.
- 6 Francis, J. G. F., The QR transformation, parts 1 and 2, Comp. J. 4 (1961), 265-271 and 332-345.
- 7 Ruhe, A., Eigenvalues of a complex matrix by the QR-method, ALGOL progr. contr. no. 19, BIT 6 (1966), 350-358.
- 8 Dekker T. J. and W. Hoffmann, ALGOL 60 procedures in numerical algebra, part 2, Mathematical Centre Tracts 23, Amsterdam 1968.
- 9 Wilkinson, J. H., Global convergence of tridiagonal QR algorithm with origin shifts (to appear in Linear algebra and its applications 1 No. 3).
- 10 Hyman, M. A., Eigenvalues and eigenvectors of general matrices, 12th National Meeting ACM, Houston, Texas 1957.
- 11 Parlett, B., Laguerre's method applied to the matrix eigenvalue problem, Math. Comp. 18 (1964), 464-485.
- 12 Dekker, T. J., Newton-Laguerre iteration, MR82, Mathematical Centre, Amsterdam 1966, and Programmation en mathématiques numériques, C.N.R.S. (Paris 1968), 189-200.
- 13 Varah, J. M., Eigenvectors of a real matrix by inverse iteration, Techn. Rep. CS34, Stanford University 1966.
- 14 Varah, J. M., The computation of bounds for the invariant subspaces of a general matrix operator, Techn. Rept. CS66, Stanford University 1967.
- 15 Businger, P. A., Eigenvalues and eigenvectors of a real symmetric matrix by the QR method, Algorithm 254, Comm. ACM 8 (1965), 218-219.
- 16 Osborne, E. E., On preconditioning of matrices, J. ACM 7 (1960), 338-345.
- 17 Eberlein, P. J., A Jacobi-like method for the automatic computation of eigenvalues and eigenvectors of an arbitrary matrix, J. SIAM 10 (1962), 74-88 and 393.
- 18 Eberlein P. J. and J. Boothroyd, Solution to the eigenproblem by a norm reducing Jacobi type method, Num. Mat. 11 (1968), 1-12.
- 19 Paardekoooper, M. H. C., Normalizing of matrices, Technological University, Eindhoven (to appear).

BOEKBESPREKINGEN

COMPUTER STUDIES IN THE HUMANITIES AND VERBAL BEHAVIOR EEN NIEUW TIJDSCHRIFT

Van de Uitgeverij Mouton & Co te Den Haag ontvingen wij het eerste nummer van het nieuwe tijdschrift, getiteld „Computer studies in the humanities and verbal behavior”. Het blad startte in januari 1968, verschijnt per kwartaal, en kost in abonnement f 36.— per jaar. Het blad wordt geschraagd door de universiteiten van Colorado, Kansas en North Carolina. De redactie wordt ter zijde gestaan door adviseurs voor anthropologie, archeologie, kunst, klassieken, opvoeding, folklore, historie, linguïstiek, literatuur, wiskunde, muziek, filosofie, politieke wetenschappen, psycholinguïstiek, sociologie, toneel, enz.

Uit deze opsomming wordt wel duidelijk, dat heel het gebied van de humaniora (om dit woord maar als vertaling te gebruiken, „menswetenschappen” heeft een bijklank) moet worden bestreken. Voor al deze „humanities” zet de tijdschrifttitel het woord computer, dus men gaat er van uit, dat al deze terreinen iets met computers van doen hebben. Hoe de schakel tussen computer en deze kennisgebieden ligt, moet uit de artikelen zelf duidelijk worden. Dat de tijdschrifttitel achter de term „humanities” nog toevoegt: „and verbal behavior”