

```

SUBROUTINE eigsrt(d,v,n,np)
INTEGER n,np
REAL d(np),v(np,np)
    Given the eigenvalues d and eigenvectors v as output from jacobi (§11.1) or tqli (§11.3),
    this routine sorts the eigenvalues into descending order, and rearranges the columns of v
    correspondingly. The method is straight insertion.
INTEGER i,j,k
REAL p
do 13 i=1,n-1
    k=i
    p=d(i)
    do 11 j=i+1,n
        if(d(j).ge.p)then
            k=j
            p=d(j)
        endif
    enddo 11
    if(k.ne.i)then
        d(k)=d(i)
        d(i)=p
        do 12 j=1,n
            p=v(j,i)
            v(j,i)=v(j,k)
            v(j,k)=p
        enddo 12
    endif
enddo 13
return
END

```

CITED REFERENCES AND FURTHER READING:

- Golub, G.H., and Van Loan, C.F. 1989, *Matrix Computations*, 2nd ed. (Baltimore: Johns Hopkins University Press), §8.4.
- Smith, B.T., et al. 1976, *Matrix Eigensystem Routines — EISPACK Guide*, 2nd ed., vol. 6 of *Lecture Notes in Computer Science* (New York: Springer-Verlag). [1]
- Wilkinson, J.H., and Reinsch, C. 1971, *Linear Algebra*, vol. II of *Handbook for Automatic Computation* (New York: Springer-Verlag). [2]

11.2 Reduction of a Symmetric Matrix to Tridiagonal Form: Givens and Householder Reductions

As already mentioned, the optimum strategy for finding eigenvalues and eigenvectors is, first, to reduce the matrix to a simple form, only then beginning an iterative procedure. For symmetric matrices, the preferred simple form is tridiagonal. The *Givens reduction* is a modification of the Jacobi method. Instead of trying to reduce the matrix all the way to diagonal form, we are content to stop when the matrix is tridiagonal. This allows the procedure to be carried out *in a finite number of steps*, unlike the Jacobi method, which requires iteration to convergence.

Givens Method

For the Givens method, we choose the rotation angle in equation (11.1.1) so as to zero an element that is *not* at one of the four “corners,” i.e., not a_{pp} , a_{pq} , or a_{qq} in equation (11.1.3). Specifically, we first choose \mathbf{P}_{23} to annihilate a_{31} (and, by symmetry, a_{13}). Then we choose \mathbf{P}_{24} to annihilate a_{41} . In general, we choose the sequence

$$\mathbf{P}_{23}, \mathbf{P}_{24}, \dots, \mathbf{P}_{2n}; \mathbf{P}_{34}, \dots, \mathbf{P}_{3n}; \dots; \mathbf{P}_{n-1,n}$$

where \mathbf{P}_{jk} annihilates $a_{k,j-1}$. The method works because elements such as a'_{rp} and a'_{rq} , with $r \neq p \neq q$, are linear combinations of the old quantities a_{rp} and a_{rq} , by equation (11.1.4). Thus, if a_{rp} and a_{rq} have already been set to zero, they remain zero as the reduction proceeds. Evidently, of order $n^2/2$ rotations are required, and the number of multiplications in a straightforward implementation is of order $4n^3/3$, not counting those for keeping track of the product of the transformation matrices, required for the eigenvectors.

The Householder method, to be discussed next, is just as stable as the Givens reduction and it is a factor of 2 more efficient, so the Givens method is not generally used. Recent work (see [1]) has shown that the Givens reduction can be reformulated to reduce the number of operations by a factor of 2, and also avoid the necessity of taking square roots. This appears to make the algorithm competitive with the Householder reduction. However, this “fast Givens” reduction has to be monitored to avoid overflows, and the variables have to be periodically rescaled. There does not seem to be any compelling reason to prefer the Givens reduction over the Householder method.

Householder Method

The Householder algorithm reduces an $n \times n$ symmetric matrix \mathbf{A} to tridiagonal form by $n - 2$ orthogonal transformations. Each transformation annihilates the required part of a whole column and whole corresponding row. The basic ingredient is a Householder matrix \mathbf{P} , which has the form

$$\mathbf{P} = \mathbf{1} - 2\mathbf{w} \cdot \mathbf{w}^T \quad (11.2.1)$$

where \mathbf{w} is a real vector with $|\mathbf{w}|^2 = 1$. (In the present notation, the *outer* or matrix product of two vectors, \mathbf{a} and \mathbf{b} is written $\mathbf{a} \cdot \mathbf{b}^T$, while the *inner* or scalar product of the vectors is written as $\mathbf{a}^T \cdot \mathbf{b}$.) The matrix \mathbf{P} is orthogonal, because

$$\begin{aligned} \mathbf{P}^2 &= (\mathbf{1} - 2\mathbf{w} \cdot \mathbf{w}^T) \cdot (\mathbf{1} - 2\mathbf{w} \cdot \mathbf{w}^T) \\ &= \mathbf{1} - 4\mathbf{w} \cdot \mathbf{w}^T + 4\mathbf{w} \cdot (\mathbf{w}^T \cdot \mathbf{w}) \cdot \mathbf{w}^T \\ &= \mathbf{1} \end{aligned} \quad (11.2.2)$$

Therefore $\mathbf{P} = \mathbf{P}^{-1}$. But $\mathbf{P}^T = \mathbf{P}$, and so $\mathbf{P}^T = \mathbf{P}^{-1}$, proving orthogonality.

Rewrite \mathbf{P} as

$$\mathbf{P} = \mathbf{I} - \frac{\mathbf{u} \cdot \mathbf{u}^T}{H} \quad (11.2.3)$$

where the scalar H is

$$H \equiv \frac{1}{2} |\mathbf{u}|^2 \quad (11.2.4)$$

and \mathbf{u} can now be any vector. Suppose \mathbf{x} is the vector composed of the first column of \mathbf{A} . Choose

$$\mathbf{u} = \mathbf{x} \mp |\mathbf{x}| \mathbf{e}_1 \quad (11.2.5)$$

where \mathbf{e}_1 is the unit vector $[1, 0, \dots, 0]^T$, and the choice of signs will be made later. Then

$$\begin{aligned} \mathbf{P} \cdot \mathbf{x} &= \mathbf{x} - \frac{\mathbf{u}}{H} \cdot (\mathbf{x} \mp |\mathbf{x}| \mathbf{e}_1)^T \cdot \mathbf{x} \\ &= \mathbf{x} - \frac{2\mathbf{u} \cdot (|\mathbf{x}|^2 \mp |\mathbf{x}|x_1)}{2|\mathbf{x}|^2 \mp 2|\mathbf{x}|x_1} \\ &= \mathbf{x} - \mathbf{u} \\ &= \pm |\mathbf{x}| \mathbf{e}_1 \end{aligned} \quad (11.2.6)$$

This shows that the Householder matrix \mathbf{P} acts on a given vector \mathbf{x} to zero all its elements except the first one.

To reduce a symmetric matrix \mathbf{A} to tridiagonal form, we choose the vector \mathbf{x} for the first Householder matrix to be the lower $n - 1$ elements of the first column. Then the lower $n - 2$ elements will be zeroed:

$$\begin{aligned} \mathbf{P}_1 \cdot \mathbf{A} &= \left[\begin{array}{c|ccc} 1 & 0 & 0 & \cdots & 0 \\ \hline 0 & & & & \\ 0 & & & & \\ \vdots & & & & \\ 0 & & & & \end{array} \right] \cdot \left[\begin{array}{c|ccc} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ \hline a_{21} & & & & \\ a_{31} & & & & \\ \vdots & & & & \\ a_{n1} & & & & \end{array} \right] \\ &= \left[\begin{array}{c|ccc} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ \hline k & & & & \\ 0 & & & & \\ \vdots & & & & \\ 0 & & & & \end{array} \right] \quad (11.2.7) \end{aligned}$$

Here we have written the matrices in partitioned form, with ${}^{(n-1)}\mathbf{P}$ denoting a Householder matrix with dimensions $(n - 1) \times (n - 1)$. The quantity k is simply plus or minus the magnitude of the vector $[a_{21}, \dots, a_{n1}]^T$.

The complete orthogonal transformation is now

$$\mathbf{A}' = \mathbf{P} \cdot \mathbf{A} \cdot \mathbf{P} = \left[\begin{array}{c|cccc} a_{11} & k & 0 & \cdots & 0 \\ \hline k & & & & \\ 0 & & & & \\ \vdots & & & & \\ 0 & & & & \end{array} \right] \quad (11.2.8)$$

irrelevant

We have used the fact that $\mathbf{P}^T = \mathbf{P}$.

Now choose the vector \mathbf{x} for the second Householder matrix to be the bottom $n - 2$ elements of the second column, and from it construct

$$\mathbf{P}_2 \equiv \left[\begin{array}{c|cccc} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \hline 0 & 0 & & & \\ \vdots & \vdots & & & \\ 0 & 0 & & & \end{array} \right] \quad (11.2.9)$$

$(n-2)\mathbf{P}_2$

The identity block in the upper left corner insures that the tridiagonalization achieved in the first step will not be spoiled by this one, while the $(n - 2)$ -dimensional Householder matrix $(n-2)\mathbf{P}_2$ creates one additional row and column of the tridiagonal output. Clearly, a sequence of $n - 2$ such transformations will reduce the matrix \mathbf{A} to tridiagonal form.

Instead of actually carrying out the matrix multiplications in $\mathbf{P} \cdot \mathbf{A} \cdot \mathbf{P}$, we compute a vector

$$\mathbf{p} \equiv \frac{\mathbf{A} \cdot \mathbf{u}}{H} \quad (11.2.10)$$

Then

$$\begin{aligned} \mathbf{A} \cdot \mathbf{P} &= \mathbf{A} \cdot \left(1 - \frac{\mathbf{u} \cdot \mathbf{u}^T}{H}\right) = \mathbf{A} - \mathbf{p} \cdot \mathbf{u}^T \\ \mathbf{A}' = \mathbf{P} \cdot \mathbf{A} \cdot \mathbf{P} &= \mathbf{A} - \mathbf{p} \cdot \mathbf{u}^T - \mathbf{u} \cdot \mathbf{p}^T + 2K\mathbf{u} \cdot \mathbf{u}^T \end{aligned}$$

where the scalar K is defined by

$$K = \frac{\mathbf{u}^T \cdot \mathbf{p}}{2H} \quad (11.2.11)$$

If we write

$$\mathbf{q} \equiv \mathbf{p} - K\mathbf{u} \quad (11.2.12)$$

then we have

$$\mathbf{A}' = \mathbf{A} - \mathbf{q} \cdot \mathbf{u}^T - \mathbf{u} \cdot \mathbf{q}^T \quad (11.2.13)$$

This is the computationally useful formula.

Following [2], the routine for Householder reduction given below actually starts in the n th column of \mathbf{A} , not the first as in the explanation above. In detail, the equations are as follows: At stage m ($m = 1, 2, \dots, n - 2$) the vector \mathbf{u} has the form

$$\mathbf{u}^T = [a_{i1}, a_{i2}, \dots, a_{i,i-2}, a_{i,i-1} \pm \sqrt{\sigma}, 0, \dots, 0] \quad (11.2.14)$$

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMs visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).

Here

$$i \equiv n - m + 1 = n, n - 1, \dots, 3 \quad (11.2.15)$$

and the quantity σ ($|x|^2$ in our earlier notation) is

$$\sigma = (a_{i1})^2 + \dots + (a_{i,i-1})^2 \quad (11.2.16)$$

We choose the sign of σ in (11.2.14) to be the same as the sign of $a_{i,i-1}$ to lessen roundoff error.

Variables are thus computed in the following order: $\sigma, \mathbf{u}, H, \mathbf{p}, K, \mathbf{q}, \mathbf{A}'$. At any stage m , \mathbf{A} is tridiagonal in its last $m - 1$ rows and columns.

If the eigenvectors of the final tridiagonal matrix are found (for example, by the routine in the next section), then the eigenvectors of \mathbf{A} can be obtained by applying the accumulated transformation

$$\mathbf{Q} = \mathbf{P}_1 \cdot \mathbf{P}_2 \cdots \mathbf{P}_{n-2} \quad (11.2.17)$$

to those eigenvectors. We therefore form \mathbf{Q} by recursion after all the \mathbf{P} 's have been determined:

$$\begin{aligned} \mathbf{Q}_{n-2} &= \mathbf{P}_{n-2} \\ \mathbf{Q}_j &= \mathbf{P}_j \cdot \mathbf{Q}_{j+1}, \quad j = n - 3, \dots, 1 \\ \mathbf{Q} &= \mathbf{Q}_1 \end{aligned} \quad (11.2.18)$$

The input parameters for the routine below are the $n \times n$ real, symmetric matrix \mathbf{a} , stored in an $np \times np$ array. On output, \mathbf{a} contains the elements of the orthogonal matrix \mathbf{q} . The vector \mathbf{d} returns the diagonal elements of the tridiagonal matrix \mathbf{A}' , while the vector \mathbf{e} returns the off-diagonal elements in its components 2 through n , with $e(1) = 0$. Note that since \mathbf{a} is overwritten, you should copy it before calling the routine, if it is required for subsequent computations.

No extra storage arrays are needed for the intermediate results. At stage m , the vectors \mathbf{p} and \mathbf{q} are nonzero only in elements $1, \dots, i$ (recall that $i = n - m + 1$), while \mathbf{u} is nonzero only in elements $1, \dots, i - 1$. The elements of the vector \mathbf{e} are being determined in the order $n, n - 1, \dots$, so we can store \mathbf{p} in the elements of \mathbf{e} not already determined. The vector \mathbf{q} can overwrite \mathbf{p} once \mathbf{p} is no longer needed. We store \mathbf{u} in the i th row of \mathbf{a} and \mathbf{u}/H in the i th column of \mathbf{a} . Once the reduction is complete, we compute the matrices \mathbf{Q}_j using the quantities \mathbf{u} and \mathbf{u}/H that have been stored in \mathbf{a} . Since \mathbf{Q}_j is an identity matrix in the last $n - j + 1$ rows and columns, we only need compute its elements up to row and column $n - j$. These can overwrite the \mathbf{u} 's and \mathbf{u}/H 's in the corresponding rows and columns of \mathbf{a} , which are no longer required for subsequent \mathbf{Q} 's.

The routine `tred2`, given below, includes one further refinement. If the quantity σ is zero or "small" at any stage, one can skip the corresponding transformation. A simple criterion, such as

$$\sigma < \frac{\text{smallest positive number representable on machine}}{\text{machine precision}}$$

would be fine most of the time. A more careful criterion is actually used. Define the quantity

$$\epsilon = \sum_{k=1}^{i-1} |a_{ik}| \quad (11.2.19)$$

If $\epsilon = 0$ to machine precision, we skip the transformation. Otherwise we redefine

$$a_{ik} \text{ becomes } a_{ik}/\epsilon \quad (11.2.20)$$

and use the scaled variables for the transformation. (A Householder transformation depends only on the ratios of the elements.)

Note that when dealing with a matrix whose elements vary over many orders of magnitude, it is important that the matrix be permuted, insofar as possible, so that the smaller elements are in the top left-hand corner. This is because the reduction is performed starting from the bottom right-hand corner, and a mixture of small and large elements there can lead to considerable rounding errors.

The routine `tred2` is designed for use with the routine `tqli` of the next section. `tqli` finds the eigenvalues and eigenvectors of a symmetric, tridiagonal matrix. The combination of `tred2` and `tqli` is the most efficient known technique for finding all the eigenvalues and eigenvectors (or just all the eigenvalues) of a real, symmetric matrix.

In the listing below, the statements indicated by comments are required only for subsequent computation of eigenvectors. If only eigenvalues are required, omission of the commented statements speeds up the execution time of `tred2` by a factor of 2 for large n . In the limit of large n , the operation count of the Householder reduction is $2n^3/3$ for eigenvalues only, and $4n^3/3$ for both eigenvalues and eigenvectors.

```

SUBROUTINE tred2(a,n,np,d,e)
INTEGER n,np
REAL a(np,np),d(np),e(np)
  Householder reduction of a real, symmetric, n by n matrix a, stored in an np by np physical
  array. On output, a is replaced by the orthogonal matrix Q effecting the transformation. d
  returns the diagonal elements of the tridiagonal matrix, and e the off-diagonal elements,
  with e(1)=0. Several statements, as noted in comments, can be omitted if only eigenvalues
  are to be found, in which case a contains no useful information on output. Otherwise they
  are to be included.
INTEGER i,j,k,l
REAL f,g,h,hh,scale
do 18 i=n,2,-1
  l=i-1
  h=0.
  scale=0.
  if(l.gt.1)then
    do 11 k=1,l
      scale=scale+abs(a(i,k))
    enddo 11
    if(scale.eq.0.)then
      e(i)=a(i,l)
    else
      do 12 k=1,l
        a(i,k)=a(i,k)/scale
        h=h+a(i,k)**2
      enddo 12
      f=a(i,l)

```

Skip transformation.

Use scaled a's for transformation.
Form σ in h.

```

    g=-sign(sqrt(h),f)
    e(i)=scale*g
    h=h-f*g
    a(i,1)=f-g
    f=0.
    do 15 j=1,1
C Omit following line if finding only eigenvalues
    a(j,i)=a(i,j)/h
    g=0.
    do 13 k=1,j
        g=g+a(j,k)*a(i,k)
    enddo 13
    do 14 k=j+1,1
        g=g+a(k,j)*a(i,k)
    enddo 14
    e(j)=g/h
    f=f+e(j)*a(i,j)
    enddo 15
    hh=f/(h+h)
    do 17 j=1,1
        f=a(i,j)
        g=e(j)-hh*f
        e(j)=g
    do 16 k=1,j
        a(j,k)=a(j,k)-f*e(k)-g*a(i,k)
    enddo 16
    enddo 17
    endif
    else
        e(i)=a(i,1)
    endif
    d(i)=h
enddo 18
C Omit following line if finding only eigenvalues.
d(1)=0.
e(1)=0.
do 24 i=1,n
    Begin accumulation of transformation matrices.
C Delete lines from here ...
    l=i-1
    if(d(i).ne.0.)then
        This block skipped when i=1.
        do 22 j=1,1
            g=0.
            do 19 k=1,1
                Use u and u/H stored in a to form P · Q.
                g=g+a(i,k)*a(k,j)
            enddo 19
            do 21 k=1,1
                a(k,j)=a(k,j)-g*a(k,i)
            enddo 21
        enddo 22
    endif
C ... to here when finding only eigenvalues.
d(i)=a(i,i)
    This statement remains.
C Also delete lines from here ...
a(i,i)=1.
    Reset row and column of a to identity matrix for
    do 23 j=1,1
        a(i,j)=0.
        a(j,i)=0.
    enddo 23
C ... to here when finding only eigenvalues.
enddo 24
return
END

```

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMs visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).

CITED REFERENCES AND FURTHER READING:

- Golub, G.H., and Van Loan, C.F. 1989, *Matrix Computations*, 2nd ed. (Baltimore: Johns Hopkins University Press), §5.1. [1]
- Smith, B.T., et al. 1976, *Matrix Eigensystem Routines — EISPACK Guide*, 2nd ed., vol. 6 of *Lecture Notes in Computer Science* (New York: Springer-Verlag).
- Wilkinson, J.H., and Reinsch, C. 1971, *Linear Algebra*, vol. II of *Handbook for Automatic Computation* (New York: Springer-Verlag). [2]

11.3 Eigenvalues and Eigenvectors of a Tridiagonal Matrix

Evaluation of the Characteristic Polynomial

Once our original, real, symmetric matrix has been reduced to tridiagonal form, one possible way to determine its eigenvalues is to find the roots of the characteristic polynomial $p_n(\lambda)$ directly. The characteristic polynomial of a tridiagonal matrix can be evaluated for any trial value of λ by an efficient recursion relation (see [1], for example). The polynomials of lower degree produced during the recurrence form a Sturmian sequence that can be used to localize the eigenvalues to intervals on the real axis. A root-finding method such as bisection or Newton's method can then be employed to refine the intervals. The corresponding eigenvectors can then be found by inverse iteration (see §11.7).

Procedures based on these ideas can be found in [2,3]. If, however, more than a small fraction of all the eigenvalues and eigenvectors are required, then the factorization method next considered is much more efficient.

The QR and QL Algorithms

The basic idea behind the *QR* algorithm is that any real matrix can be decomposed in the form

$$\mathbf{A} = \mathbf{Q} \cdot \mathbf{R} \quad (11.3.1)$$

where \mathbf{Q} is orthogonal and \mathbf{R} is upper triangular. For a general matrix, the decomposition is constructed by applying Householder transformations to annihilate successive columns of \mathbf{A} below the diagonal (see §2.10).

Now consider the matrix formed by writing the factors in (11.3.1) in the opposite order:

$$\mathbf{A}' = \mathbf{R} \cdot \mathbf{Q} \quad (11.3.2)$$

Since \mathbf{Q} is orthogonal, equation (11.3.1) gives $\mathbf{R} = \mathbf{Q}^T \cdot \mathbf{A}$. Thus equation (11.3.2) becomes

$$\mathbf{A}' = \mathbf{Q}^T \cdot \mathbf{A} \cdot \mathbf{Q} \quad (11.3.3)$$