

Chapter 5

Dense Matrix Factorisation Operations

The following routines are described in the following pages:

BKpfactor, BKPsolve	Bunch–Kaufman–Parlett factor and solve (stable symmetric factorisation)	71
CHfactor, MCHfactor, CHsolve LDLfactor, LDLsolve	Cholesky, LDL^T factor and solve	73
LUfactor, LUsolve, LUTsolve LUcondest, m_inverse	LU factor (Gaussian elimination) and solve	75
QRfactor, QRsolve, QRTsolve QRCPfactor, QRCPsolve QRcondest	QR factor and solve with/out column pivoting	77
makeQ, makeR	extract matrices from compact form (QR only)	79
givens, rot_cols rot_rows, rot_vec	compute and apply Givens' rotations	80
hhvec, hhtrvec hhtrrows, hhtrcols	Householder transformations	82
Dsolve, Lsolve, LTsolve Usolve, UTsolve	solve for diagonal and triangular matrices	84
LDLupdate, QRupdate	update routines for LDL^T and QR factorisations	86
schur, symmeig, trieig	eigenvalue routines	88
schur_evals, schur_vecs	eigenvalue/vector extraction routines	90
svd, bisvd	singular value decomposition	91
fft, ifft	Fast Fourier Transform	93

To use these routines use the include statements

```
#include "matrix.h"  
#include "matrix2.h"
```

NAME

BKPFactor, BKPSolve – Bunch–Kaufman–Parlett symmetric indefinite factorise and solve

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
MAT    *BKPFactor(A,pivot,blocks)
MAT    *A;
PERM   *pivot, *blocks;

VEC    *BKPSolve(A,pivot,blocks,b,x)
MAT    *A;
PERM   *pivot, *blocks;
VEC    *b, *x;
```

DESCRIPTION

The routine BKPFactor() forms *in situ* a symmetric indefinite factorisation of the matrix **A** of the form

$$P^T A P = M D M^T$$

where P is a permutation matrix, M is lower triangular, and D is block diagonal, with 1×1 or 2×2 blocks. The matrix P is represented by the permutation **pivot** and D_{ii} is a 1×1 block if and only if **blocks->pe[i] == i**; otherwise **blocks->pe[i]** is the index of the other row/column in the 2×2 block. After the routine the D and M factors are stored in **A** in compact form. This avoids the requirement for additional vectors or matrices for storage.

Note that **pivot** and **blocks** must both be non-NULL and different for both BKPFactor() and BKPSolve().

The routine BKPSolve() solves the equation $Ax = b$ for x . The solve routine BKPSolve() is designed specifically to work with BKPFactor() as they operate on the same compact storage scheme. Note that the factorisation may succeed when the matrix **A** passed is singular, and that the solve routine may then fail, raising an E_SING error. The solve routine may be used *in situ* with **b == x**. If **x** is NULL or too small to hold the result, then a new vector is created of the appropriate size for storing the result. In either case the resulting solution vector is returned.

This factorisation routine, and the accompanying solve routine are derived from “Decomposition of a Symmetric Matrix” by J. Bunch, L. Kaufman and B. Parlett, *Numerische Mathematik* **27**, 95–109 (1976).

Errors will be raised if **A** or **pivot** or **blocks** are NULL, or if **A** is not square, or if the sizes of **A**, **pivot** or **blocks** are not compatible.

EXAMPLE

```
MAT    *A;
PERM   *pivot, *blocks;
VEC    *x, *b;
.....
A = in_mat(MNULL);
b = in_vec(VNULL);
pivot = get_perm(A->m);
blocks = get_perm(A->m);
/* assuming A symmetric */
BKPFactor(A,pivot,blocks);
x = BKPSolve(A,pivot,blocks,b,VNULL);
```

SEE ALSO

`CHfactor()` and `CHsolve()`

SOURCE FILE: `BKPfactor.c`

NAME

CHfactor, MCHfactor, CHsolve, LDLfactor, LDLsolve – Cholesky factor and solve

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
MAT    *CHfactor(A)
MAT    *A;

MAT    *MCHfactor(A,tol)
MAT    *A;
double tol;

VEC    *CHsolve(A,b,x)
MAT    *A;
VEC    *b, *x;

MAT    *LDLfactor(A)
MAT    *A;

VEC    *LDLsolve(A,b,x)
MAT    *A;
VEC    *b, *x;
```

DESCRIPTION

Both `CHfactor()` and `LDLfactor()` factor the matrix A *in situ* and returns the factored matrix (in compact form). The Cholesky factorisation routine and the LDL^T routines both use only the lower triangular part of A , but the Cholesky factorisation routine fills the upper triangular part of A also.

These routines require that A is square. The Cholesky factorisation, in particular, requires that A be sufficiently positive definite (e.g. lowest eigenvalue of A is at least machine epsilon away from zero). If non-positive definiteness is detected during factorisation, then an `E_POSDEF` error will be raised. If you wish to catch such an error, see information on the `catch()` macro. If your matrix is indefinite, then it would be best to use the `BKPFactor()` and `BKPSolve()` routines.

The routine `MCHfactor()` computes a *modified* Cholesky factorisation. This is not a true Cholesky factorisation, but rather the Cholesky factorisation of $A + D$ where D is a diagonal matrix with non-negative entries. Whether the A matrix is modified in this way is determined by the `tol` parameter; the diagonal entry of the Cholesky factorisation is ensured to be $\geq \sqrt{\text{tol}}$. The D matrix is guaranteed to be zero in exact arithmetic if $u^T A u \geq \text{tol} u^T u$ for all u .

EXAMPLE

```
MAT    *A, *LLT, *LDL;
VEC    *b, *x;
double tol;
.....
A = in_mat(MNULL);
b = in_vec(VNULL);
input("Input tol for modified Cholesky: ", "%lf", &tol);
LLT = cp_mat(A,MNULL);
```

```
/* If A positive definite... */
CHfactor(LLT);
x = CHsolve(LLT,b,VNULL);
/* ...otherwise, get approximate solution... */
LLT = cp_mat(A,MNULL);
MCHfactor(LLT,tol);      /* LLT now has factors of A + D */
MCHsolve(LLT,b,x);
/* ...or use LDL factorisation */
LDL = cp_mat(A,MNULL);
LDLfactor(LDL);
LDLsolve(LDL,b,x);
```

SEE ALSO

catch() and BKPfactor()

SOURCE FILE: CHfactor.c

NAME

LUfactor, LUsolve, LUTsolve, LUcondest, m_inverse – *LU* factorisation (Gaussian elimination) and solve

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
MAT    *LUfactor(A,pivot)
MAT    *A;
PERM   *pivot;

VEC    *LUsolve(A,pivot,b,x)
MAT    *A;
PERM   *pivot;
VEC    *b, *x;

VEC    *LUTsolve(A,pivot,b,x)
MAT    *A;
PERM   *pivot;
VEC    *b, *x;

double LUcondest(LU,pivot)
MAT    *LU;
PERM   *pivot;

MAT    *m_inverse(A,out)
MAT    *A, *out;
```

DESCRIPTION

The routine `LUfactor()` performs *LU* factorisation, which is otherwise known as Gaussian elimination with implicit scaled partial pivoting. The *LU* factors of *A* are stored in *A* in compact form. Once this is done, the routine `LUsolve()` can be used to solve equations of the form $Ax = b$ for x by forward and back substitution. The system $A^T x = b$ can be solved by using `LUTsolve()`. The code for a full factorise and solving $Ax = b$ and $A^T y = b$ is:

```
/* set up A and b */
.....
pivot = get_perm(A->m);
x = get_vec(A->n);
y = get_vec(A->m);
LU = cp_mat(A,MNULL);
LUfactor(LU,pivot);
x = LUsolve(LU,pivot,b,x);
y = LUTsolve(LU,pivot,b,y);
condition = LUcondest(LU,pivot);
```

A full description of Gaussian elimination with partial pivoting and its numerical behaviour can be found in a number of books, though we refer the reader specifically to *Matrix Computations* by G.H. Golub and C. van Loan, North Oxford Academic, §§3.2–3.4, pp. 92–122, 2nd Edition (1989). The variant here is that scaling is used *implicitly*. That is, scaling is only used to decide which rows to swap during the partial pivoting process.

Note that the factorisation routine `LUfactor()` may succeed where the solve routine `LUsolve()` fails if, for example, `A` is singular. Also note that *LU* factorisation also succeeds when `A` is not even square, though this is a requirement for the success of `LUsolve()`. Errors are raised by `LUfactor()` if `A` or `pivot` is `NULL`, if the size of `pivot` is less than the number of rows of `A`. Errors are raised by `LUsolve()` if these conditions occur, if `b` is `NULL`, if `A` is not square. Then if `x` is `NULL` or too small to contain the result a new vector of the appropriate size is created. In either case the solution of $Ax = b$, `x`, is returned. The routines `LUsolve()` and `LUTsolve()` may *not* be used *in situ*; that is, with `b == x`. This is because of the permutation of the vector.

The condition number (relative to the infinity norm) can be *estimated* using the routine `LUcondest()`. This estimate is not guaranteed to under- or over-estimate the true condition number; however, it can usually be relied on to give an estimate correct to within an order of magnitude, which is usually all that is required.

The routine `m_inverse()` computes the inverse of `A` and returns it in `out`. This is carried out using the *LU* factorisation routines. As is usually noted in numerical analysis texts, inverse matrices should rarely be computed. If a system of equations need to be solved, use the above code calling `LUfactor()` and `LUsolve()` directly.

SOURCE FILE: `LUfactor.c`

NAME

QRfactor, QRCPfactor, QRsolve, QRTsolve, QRcondest – QR factorisation and solve

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
MAT    *QRfactor(A,diag)
MAT    *A;
VEC    *diag;

MAT    *QRCPfactor(A,diag,pivot)
MAT    *A;
VEC    *diag;
PERM   *pivot;

VEC    *QRsolve(A,diag,b,x)
MAT    *A;
VEC    *diag;
VEC    *b, *x;

VEC    *QRTsolve(A,diag,b,x)
MAT    *A;
VEC    *diag;
VEC    *b, *x;

VEC    *QRCPsolve(A,diag,pivot,b,x)
MAT    *A;
VEC    *diag;
VEC    *b, *x;

double QRcondest(QR)
MAT    *QR;
```

DESCRIPTION

The routine `QRfactor()` performs a straightforward QR factorisation of A . For those unfamiliar with the terminology, the QR factorisation of A is a factorisation of the form

$$A = QR$$

where Q is orthogonal and R is upper triangular. This factorisation exists whether or not A is singular or even square. The QR factorisation is performed using Householder transformations. (These are orthogonal matrices of the form $P_i = I - \alpha_i v_i v_i^T$ where $\alpha_i = 2/v_i^T v_i$.)

The routine `QRCPfactor()` performs a QR factorisation with column pivoting, which is a factorisation of the form

$$A\Pi^T = QR$$

where additionally, Π is a permutation matrix. The Π matrix is represented by `pivot`. This is done exactly as for `QRfactor()` except for the pivoting.

Both of these factorisations are performed *in situ*, and store the Q and R factors compactly in A and `diag`. This compact form is used consistently within this package, and is essentially that of Golub and

van Loan's *Matrix Computations*, §5.2, p. 212, 2nd edition, (1989) except that the v 's are not normalised in this package. The dimensions of both `diag` must be at least as large as the minimum of the number of rows and columns of A .

Once A , `diag` contain this compact representation of the QR factors of A , we can use `QRsolve()` to solve systems of linear equations, and indeed, find least square error solutions to overdetermined systems of equations. See *Matrix Computations*, §1.4, p. 11 for an example. Indeed, the code

```
MAT      *QR;
        .....
QR = cp_mat(A,MNULL);
QRfactor(QR,diag);
QRsolve(QR,diag,b,x);
```

finds the least squares solution x to

$$Ax \approx b.$$

Similarly, if `QRCPfactor()` is to be used to factor A , then `QRCPsolve()` can be used to solve the least squares problem $Ax \approx b$. The code to do this is:

```
QR = cp_mat(A,MNULL);
QRCPfactor(QR,diag,pivot);
QRCPsolve(QR,diag,pivot,b,x);
```

Note that `QRtsolve(QR,diag,b,x)` solves the *underdetermined* problem $Ax = b$; that is, it computes the minimum 2-norm x that satisfies $Ax = b$ for $m \leq n$.

The condition number of a matrix factored using either `QRfactor()` or `QRCPfactor()` can be estimated using `QRconddest()`:

```
printf("2-norm condition number approx. = %g\n", QRconddest(QR));
```

The function `QRconddest()` returns a *lower bound* for the least squares condition number of the factored matrix A

$$\kappa_{LS}(A) = \|A\|_2 \|A^+\|_2$$

provided A has full rank. If A is square, then this is exactly equal to the 2-norm condition number

$$\kappa_2(A) = \|A\|_2 \|A^{-1}\|_2.$$

If the QR factors are exactly singular, then `QRconddest()` will return HUGE (HUGE_VAL for ANSI C).

The estimate is obtained by obtaining estimates for $\|R\|_2$ and $\|R^{-1}\|_2$. Note that Q and Π do not affect the 2-norm or least squares condition numbers. The estimate of $\|R^{-1}\|_2$ is found using the techniques of `LUconddest()` to obtain a vector y with unit ∞ -norm such that $\|R^{-1}y\|_\infty$ is quite small. This is described in Golub and van Loan, 2nd Edition pp. 128–130, (1989). Then the power method is applied to the matrix $(R^T R)^{-1}$ a total of three times with initial vector y . The corresponding estimate of $\|R\|_2$ is obtained by a related method of finding a vector y with unit ∞ -norm and $\|Ry\|_\infty$ quite large. The power method is applied to the matrix $R^T R$. Taking square root of the estimated eigenvalues gives a lower bound to the 2-norm condition number of R .

A simple, and usually reliable, estimate of the rank of a matrix is to factor the matrix A using `QRCPfactor()`, and then to count the number of diagonal entries of A greater than a certain tolerance in magnitude. A more reliable approach is to use the Singular Value Decomposition. See `svd()`.

SEE ALSO

Householder routines `hhvec()`, `hhtrvec()`, `hhtrrows()` and `hhtrcols()`; `svd()`.

SOURCE FILE: `QRfactor.c`

NAME

makeQ, makeR – explicitly form Q and R factors

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
MAT      *makeQ(QR,diag,Qout)
MAT      *QR;
VEC      *diag;
MAT      *Qout;

MAT      *makeR(QR,Rout)
MAT      *QR, *Rout;
```

DESCRIPTION

The routine `makeQ()` explicitly forms the orthogonal Q matrix of the QR factorisation from the compact representation in `QR`, `diag`. The result is stored in `Qout`. This routine may not be used to form `Qout` *in situ*.

The routine `makeR()` explicitly forms the upper triangular R matrix of the QR factorisation. The result is stored in `Rout`. This routine may be used *in situ*; that is, with `QR == Rout`. (Actually the routine just zeros the strictly lower triangular half of `QR`.)

If `Qout` or `Rout` is `NULL` or too small to contain the result then a new matrix is created and returned.

EXAMPLE

```
MAT      *A, *QR, *Q, *R;
VEC      *diag;
.....
diag = get_vec(A->m);
QR = cp_mat(A,MNULL);
QRfactor(QR,diag);
Q = makeQ(QR,diag,MNULL);
R = makeR(QR,MNULL);
/* makeR(QR,QR); replaces QR with the R matrix */
```

SOURCE FILE: QRfactor.c

Note that $J_{ik}(c, s)^T = J_{ik}(c, -s)$. This makes pre- and post-multiplying by transposes of $J_{ik}(c, s)$ easy.

EXAMPLE

```
int    i, k;
VEC    *x;
MAT    *A;
double c, s;
.....
/* get Givens transformation */
givens(x->ve[i],x->ve[k],&c,&s);
/* apply to x */
rot_vec(x,i,k,c,s);
/* apply symmetrically to A */
rot_cols(A,i,k,c,s);
rot_rows(A,i,k,c,s);
```

BUGS

The `givens()` routine may result in overflow if the `x` and/or `y` parameters are of size greater than $\sqrt{\text{HUGE}}$.

SOURCE FILE: `givens.c`

NAME

hhvec, hhtrcols, hhtrrows, hhtrvec – Householder transformation operations

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
VEC    *hhvec(x,i0,beta,out,newval)
VEC    *x, *out;
unsigned i0;
double *beta, *newval;

MAT    *hhtrcols(A,i0,j0,hh,beta)
MAT    *A;
unsigned i0, j0;
VEC    *hh;
double beta;

MAT    *hhtrrows(A,i0,j0,hh,beta)
MAT    *A;
unsigned i0, j0;
VEC    *hh;
double beta;

VEC    *hhtrvec(hh,beta,i0,x,out)
VEC    *hh;
double beta;
unsigned i0;
VEC    *x, *out;
```

DESCRIPTION

The routine `hhvec()` computes the parameters for a Householder transformation. In particular, given a vector x , a vector v (`== out`) and two real numbers β (`== beta`) and $newval$ where the Householder transformation $P = I - \beta vv^T$ satisfies

$$(5.3) \quad Px = \begin{bmatrix} newval \\ 0 \end{bmatrix}.$$

The `x` parameter is not modified. The formulae used are taken from *Matrix Computations* by G. Golub and C. van Loan, p. 40, 1st Edition, (1983), §5.1, pp. 196–196, 2nd Edition, (1989).

If `out` is NULL or too small to hold the v vector, then a new vector is created to store the result. In either case, the result is returned. An error are raised if the `x` vector is NULL.

The routine `hhtrcols()` forms the product AP^T where P is the Householder transformation defined by hh and β (`== beta`). (That is, $P = I - \beta hh hh^T$.) All rows i with $i < i0$ and columns j with $j < j0$ are ignored. The operations is performed *in situ* in `A`.

The routine `hhtrrows()` forms the product PA where P is the Householder transformation defined by hh and β . Again, all rows i with $i < i0$ and columns j with $j < j0$ are ignored. The operations is performed *in situ* in `A`.

Finally, the routine `hhtrvec()` forms the vector Px where P is the Householder transformation defined by hh and β . The result is stored in `out`. If `out` is NULL or too small to hold the results of the operation, then a new vector is created of the appropriate size. In either case the result is returned.

SOURCE FILE: hsehldr.c

NAME

Dsolve, Lsolve, LTsolve, Usolve, UTsolve – basic solve routines

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
VEC      *Dsolve(A,b,x)
MAT      *A;
VEC      *b, *x;

VEC      *Lsolve(A,b,x,diag)
MAT      *A;
VEC      *b, *x;
double   diag;

VEC      *LTsolve(A,b,x,diag)
MAT      *A;
VEC      *b, *x;
double   diag;

VEC      *Usolve(A,b,x,diag)
MAT      *A;
VEC      *b, *x;
double   diag;

VEC      *UTsolve(A,b,x,diag)
MAT      *A;
VEC      *b, *x;
double   diag;
```

DESCRIPTION

The routine `Dsolve()` finds and returns the solution x of $Dx = b$ where D is the diagonal part of the matrix A ($== A$).

The routine `Lsolve()` finds and returns the solution x of $Lx = b$ where L is the lower triangular part of A if `diag` is zero; L is the *strictly* lower triangular part of A with `diag` on the diagonal if `diag` is not zero. This routine uses forward substitution.

The routine `LTsolve()` finds and returns the solutions x of $L^T x = b$ where L is the lower triangular part of A if `diag` is zero; L is the *strictly* upper triangular part of A with `diag` on the diagonal if `diag` is not zero.

The routine `Usolve()` finds and returns the solution x of $Ux = b$ where U is the upper triangular part of A if `diag` is zero; U is the *strictly* upper triangular part of A with `diag` on the diagonal if `diag` is not zero. This routine uses back substitution.

The routine `UTsolve()` finds and returns the solution x of $U^T x = b$ where U is the upper triangular part of A if `diag` is zero; U is the *strictly* upper triangular part of A with `diag` on the diagonal if `diag` is not zero. This routine uses back substitution.

All of these routines may be used *in situ*; that is, they can be used with `b == x`.

If `x` is too small to contain the result then a new vector is created of the appropriate dimension. In either case the solution of the equations is returned.

The rationale behind the use of the `diag` parameter is that often, as in LU factorisation or LDL^T factorisation, the diagonal entry for L is implicit (usually one). The `diag` parameter enables these routines to be used generally, including for the results of QR factorisation, for example.

EXAMPLE

For solving $Ax = b$ using Cholesky factorisation, with only L :

```
MAT    *L;
VEC    *b, *x;
.....
Lsolve(L,b,x,0.0); /* use L's diagonal entries */
LTsolve(L,x,x,0.0);
```

For solving $Ax = b$ using LU factorisation with L unit lower triangular and no pivoting:

```
MAT    *L, *U;
VEC    *b, *x;
.....
Lsolve(L,b,x,1.0); /* L unit lower triangular */
Usolve(U,b,x,0.0);
```

SEE ALSO

`LUsolve()`, `CHsolve()`, `LDLsolve()`, `QRsolve()`

SOURCE FILE: `solve.c`

NAME

LDLupdate, QRupdate – factorisation update routines

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
MAT    *LDLupdate(LDL,w,alpha)
MAT    *LDL;
VEC    *w;
double alpha;

MAT    *QRupdate(Q,R,u,v)
MAT    *Q, *R;
VEC    *u, *v;
```

DESCRIPTION

The routine `LDLupdate()` modifies the matrix `LDL` which is assumed to contain (in compact form) the LDL^T factorisation of a matrix A . The L matrix is the strictly lower triangular part of `LDL`, except with ones on the diagonal, while D is the diagonal of `LDL`, so that $A = LDL^T$. The matrix `LDL` is modified *in situ* so that if L_+ and D_+ denote the factors described by `LDL` after the routine, then

$$L_+D_+L_+^T = A + \alpha ww^T$$

where α is the value of `alpha` and w is `w`. The modified `LDL` matrix is returned.

The method used for updating the factorisation is given in “Methods for modifying matrix factorisations” by P. Gill, G. Golub, W. Murray and M. Saunders, *Mathematics of Computations*, **28**, pp. 505–535 (1974). The particular algorithm used is algorithm C1 of their paper.

This routine may fail if $A + \alpha ww^T$ is not sufficiently positive definite; if this failure occurs, then an `E_POSDEF` error is raised.

The routine `QRupdate()` updates the QR factorisation of a matrix $A = QR$. Unlike the previous routine, this routine requires the explicit factors Q and R of A . These can be obtained from the compact form by means of the routines `makeQ()` and `makeR()`. If the matrices `Q` and `R` after the routine are denoted Q_+ and R_+ respectively, then

$$Q_+R_+ = Q(R + uv^T) = A + (Qu)v^T.$$

Setting $u = Q^T w$ gives $Q_+R_+ = A + ww^T$.

If `Q` is `NULL`, then only the `R` matrix is modified. The `R` matrix is returned.

The routine is based on one given in *Matrix Computations* by G. Golub and C. van Loan, pp. 437–443, 1st Edition (1983), pp. 593–594, 2nd Edition (1989).

EXAMPLE

Updating LDL^T factorisation:

```
MAT    *A, *LDL;
VEC    *u;
double alpha;
.....
LDL = cp_mat(A,MNULL);
LDLfactor(LDL);
.....
```

```
/* A <- A + alpha.u.u^T */  
LDLupdate(LDL,u,alpha);
```

Updating QR factorisation:

```
MAT    *A, *QR, *Q, *R;  
VEC    *diag, *beta, *u, *v, *w;  
.....  
QR = cp_mat(A,MNULL);  
QRfactor(QR,diag,beta);  
Q = makeQ(QR,diag,beta,MNULL);  
R = makeR(QR,MNULL);  
.....  
/* A <- A + w.v^T */  
u = get_vec(Q->m);  
u = vm_mlt(Q,w,u);  
QRupdate(Q,R,u,v);
```

SOURCE FILE: update.c

NAME

`schur`, `symmeig`, `trieig` – eigenvalue routines

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
MAT      *schur(A,Q)
MAT      *A, *Q;

VEC      *symmeig(A,Q,out)
MAT      *A, *Q;
VEC      *out;

VEC      *trieig(a,b,Q)
VEC      *a, *b;
MAT      *Q;
```

DESCRIPTION

The routine `schur()` computes the Real Schur decomposition of the matrix A . That is, it computes a block upper triangular matrix T and an orthogonal matrix Q such that

$$Q^T A Q = T.$$

The matrix T has diagonal blocks of sizes 1×1 and 2×2 . The eigenvalues of these diagonal blocks are the eigenvalues of the original A matrix. The algorithm used to find the eigenvalues of A is the Francis QR algorithm. This algorithm is described in *Matrix Computations* by G. Golub and C. van Loan, pp. 231–236, 1st Edition (1983), pp. 377–381, 2nd Edition (1989).

The matrix A is overwritten with T , and if Q is not `NULL` and the correct size, then the Q matrix is stored in it.

The routine `symmeig()` computes the eigenvalues of a *symmetric* matrix. It also computes an orthogonal matrix Q such that

$$Q^T A Q = \Lambda$$

where Λ is the diagonal matrix of eigenvalues. The algorithm used to find the eigenvalues of A consists of conversion to symmetric Hessenberg (symmetric tridiagonal) form, and then applying `trieig()` to obtain the eigenvalues of the tridiagonal matrix.

The eigenvalues are stored in `out` provided it is not `NULL` and is sufficiently large to contain all the eigenvalues. The vector containing the eigenvalues is returned. The matrix A is *not* overwritten.

The routine `trieig()` computes the eigenvalues of the symmetric tridiagonal matrix

$$(5.4) \quad T = \begin{bmatrix} a_0 & b_0 & & & & \\ b_0 & a_1 & b_1 & & & \\ & b_1 & a_2 & \ddots & & \\ & & \ddots & \ddots & b_{n-2} & \\ & & & b_{n-2} & a_{n-1} & \end{bmatrix}.$$

The algorithm used is a “chasing” technique described in *Matrix Computations*, pp. 278–281, 1st Edition, pp. 421–424, 2nd Edition. It also accumulates the matrix Q such that $Q^T T Q$ is diagonal. To compute the correct Q matrix, Q should be initialised to the identity matrix on entry to `trieig()`. (See `id_mat()`.)

The values in the *a* and *b* vectors are overwritten. At the end of the routine, *a* contains the eigenvalues, and the *b* vector is zero.

In all of the above routines, if the matrix *Q* is NULL on entry, then no calculation of the *Q* matrices is performed. This should speed up the routines somewhat if only the eigenvalues are needed.

EXAMPLE

Computing real Schur decomposition of (possibly) nonsymmetric *A*:

```
MAT  *A, *S, *Q, *X_re, *X_im;
VEC  *evals_re, *evals_im;
.....
S = cp_mat(A,MNULL);
Q = get_mat(A->m,A->m);
schur(S,Q);
/* get eigenvalues (real, imaginary parts) */
evals_re = get_vec(A->m);
evals_im = get_vec(A->m);
schur_evals(S,evals_re,evals_im);
/* get eigenvectors (real, imaginary parts) */
X_re = get_mat(A->m,A->m);
X_im = get_mat(A->m,A->m);
schur_evecs(S,Q,X_re,X_im);
```

Computing eigenvalues and eigenvectors of a real symmetric matrix:

```
MAT  *A, *Q;
VEC  *evals;
.....
evals = get_vec(A->m);
evals = symmeig(A,Q,evals);
```

The *Q* matrix contains the eigenvectors.

Computing the eigenvalues and eigenvectors of a symmetric tridiagonal matrix defined by the vectors *a* (the diagonal entries) and *b* (the off-diagonal entries):

```
MAT  *Q;
VEC  *a, *b;
.....
Q = get_mat(a->dim,a->dim);
id_mat(Q); /* must initialise Q */
trieig(a,b,Q);
/* a is now the vector of eigenvalues */
```

SEE ALSO

The Hessenberg routines in `hessen.c`.

BUGS

It is up to the caller of `symmeig()` to ensure that the *A* matrix is symmetric. Symmetry of *A* is neither checked nor enforced in `symmeig()`.

SOURCE FILE: `symmeig.c`, `schur.c`

NAME

`schur_evals`, `schur_vecs` – Extracting eigenvalues and eigenvectors from the Schur form

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
void schur_evals(T,re_evals,im_evals)
MAT *T;
VEC *re_evals, *im_evals;

MAT *schur_vecs(T,Q,X_re,X_im)
MAT *T, *Q;
MAT *X_re, *X_im;
```

DESCRIPTION

Both of these routines assume that T is the matrix computed by the `schur()` routine; Q is the orthogonal matrix computed by `schur()`.

The routine `schur_evals()` compute the eigenvalues of a matrix T in Schur form (block diagonal with 1×1 or 2×2 blocks). The k th eigenvalue of $A = QTQ^T$ is `re_evals->ve[k] + iim_evals->ve[k]`. At worst this requires solving a series of quadratics; however, it does simplify the task of computing eigenvalues. Complex eigenvalues come in complex conjugate pairs.

The routine `schur_vecs()` computes the matrix $X = X_re + i X_im$ such that $X^{-1}AX$ is the diagonal matrix of eigenvalues where $T = Q^T A Q$ as computed by the `schur()` routine. The columns of X are computed by means of one step of inverse iteration using the eigenvalues as computed from the Schur form. This method is usually accurate provided the eigenvalues are not too close together. The computed k th column of X is real if the computed k th eigenvalue is real. The ordering of the columns is consistent with the ordering of the eigenvalues generated by `schur_evals()`.

EXAMPLE

See example for `schur()` above.

BUGS

It is a bit difficult to check that the computed X is correct if it is complex.

SEE ALSO

`schur()`

SOURCE FILE: `schur.c`

NAME

svd, bisvd – Singular Value Decomposition routines

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
VEC      *svd(A,U,V,out)
MAT      *A, *U, *V;
VEC      *out;

VEC      *bisvd(d,f,U,V)
VEC      *d, *f;
MAT      *U, *V;
```

DESCRIPTION

The routine `svd()` performs a complete Singular Value Decomposition (SVD) on the matrix A . That is, it computes orthogonal matrices U and V such that UAV^T is diagonal and the diagonal entries are called the *singular values* of the matrix A . The first $\min(m, n)$ singular values are stored in the `out` matrix which is also returned. Note that the SVD is defined for nonsquare as well as square matrices.

If NULLs are passed for either or both U and V , then that orthogonal matrix will not be accumulated. This saves both time and space, if just the singular values are desired and not the U or V matrices. If `out` is NULL on entry to `svd()`, then a vector of the appropriate size is created to store the singular values, which is returned.

The SVD is computed by first transforming the matrix into a bidiagonal matrix (c.f. `schur()` where a matrix is transformed into Hessenberg form for eigenvalue calculations) and then applying `bisvd()`. If a matrix is already in bidiagonal form, then `bisvd()` can be called directly. The vector `d` contains the diagonal entries and `f` contains the super-diagonal entries. As for `svd()`, if NULLs are passed for either or both U and V , then that (or both) orthogonal matrix will not be accumulated. For correct results using `bisvd()`, you should initialise U and V to be identity matrices using `id_mat()` before calling `bisvd()`.

The rank of a matrix can be estimated by counting the number of singular values whose magnitude exceeds a specified tolerance. This tolerance for accurately computed matrices should probably be about 100 times `MACHEPS`; otherwise it should be about an order of magnitude larger than the errors in the matrix.

The algorithm used follows *Matrix Computations* by Golub and van Loan, pp. 430–435, 2nd Edition (1989).

EXAMPLE

For computing the SVD of A :

```
MAT      *A, *U, *V;
VEC      *svdvals;
.....
U = get_mat(A->m,A->m);
V = get_mat(A->n,A->n);
svdvals = svd(A,U,V,VNULL);
```

For computing the SVD of the bidiagonal matrix defined by d (the diagonal entries) and f (the super-diagonal entries):

```
MAT      *U, *V;
```

```
VEC    *d, *f;
.....
U = get_mat(d->dim,d->dim);
V = get_mat(d->dim,d->dim);
id_mat(U);    /* must initialise U and V */
id_mat(V);
bisvd(d,f,U,V)
/* d now contains the singular values */
```

SOURCE FILE: svd.c

NAME

fft, ifft – Fast Fourier Transform and inverse

SYNOPSIS

```
#include "matrix.h"
#include "matrix2.h"
void    fft(x_re, x_im)
VEC     *x_re, *x_im;

void    ifft(x_re, x_im)
VEC     *x_re, *x_im;
```

DESCRIPTION

The routine `fft()` performs a fast Fourier transform on the vector $x = x_re + ix_im$. The transform is computed *in situ*. It does require that the dimension of x is a power of two.

The routine `ifft()` performs the inverse fast Fourier transform of $x = x_re + ix_im$. As with `fft()` it is computed *in situ*, and the dimension of x must be a power of two.

SOURCE FILE: `fft.c`

Contents

5	Dense Matrix Factorisation Operations	70
----------	--	-----------