

NAG Library Chapter Introduction

F01 – Matrix Operations, Including Inversion

Contents

1	Scope of the Chapter	2
2	Background to the Problems	2
2.1	Matrix Inversion	2
2.2	Matrix Factorizations	3
2.3	Matrix Arithmetic and Manipulation	3
2.4	Matrix Functions	4
3	Recommendations on Choice and Use of Available Routines	4
3.1	Matrix Inversion	4
3.2	Matrix Factorizations	5
3.3	Matrix Arithmetic and Manipulation	5
3.3.1	NAG Names and LAPACK Names	6
3.4	Matrix Functions	6
4	Decision Trees	7
5	Functionality Index	10
6	Auxiliary Routines Associated with Library Routine Parameters	12
7	Routines Withdrawn or Scheduled for Withdrawal	12
8	References	13

1 Scope of the Chapter

This chapter provides facilities for four types of problem:

- (i) Matrix Inversion
- (ii) Matrix Factorizations
- (iii) Matrix Arithmetic and Manipulation
- (iv) Matrix Functions

These problems are discussed separately in Section 2.1, Section 2.2, Section 2.3 and Section 2.4.

2 Background to the Problems

2.1 Matrix Inversion

- (i) Nonsingular square matrices of order n .

If A , a square matrix of order n , is nonsingular (has rank n), then its inverse X exists and satisfies the equations $AX = XA = I$ (the identity or unit matrix).

It is worth noting that if $AX - I = R$, so that R is the ‘residual’ matrix, then a bound on the relative error is given by $\|R\|$, i.e.,

$$\frac{\|X - A^{-1}\|}{\|A^{-1}\|} \leq \|R\|.$$

- (ii) General real rectangular matrices.

A real matrix A has no inverse if it is square (n by n) and singular (has rank $< n$), or if it is of shape (m by n) with $m \neq n$, but there is a **Generalized or Pseudo-inverse** A^+ which satisfies the equations

$$AA^+A = A, \quad A^+AA^+ = A^+, \quad (AA^+)^T = AA^+, \quad (A^+A)^T = A^+A$$

(which of course are also satisfied by the inverse X of A if A is square and nonsingular).

- (a) if $m \geq n$ and $\text{rank}(A) = n$ then A can be factorized using a **QR factorization**, given by

$$A = Q \begin{pmatrix} R \\ 0 \end{pmatrix},$$

where Q is an m by m orthogonal matrix and R is an n by n , nonsingular, upper triangular matrix. The pseudo-inverse of A is then given by

$$A^+ = R^{-1}\tilde{Q}^T,$$

where \tilde{Q} consists of the first n columns of Q .

- (b) if $m \leq n$ and $\text{rank}(A) = m$ then A can be factorized using an **RQ factorization**, given by

$$A = (R \quad 0)Q^T$$

where Q is an n by n orthogonal matrix and R is an m by m , nonsingular, upper triangular matrix. The pseudo-inverse of A is then given by

$$A^+ = \tilde{Q}R^{-1},$$

where \tilde{Q} consists of the first m columns of Q .

- (c) if $m \geq n$ and $\text{rank}(A) = r \leq n$ then A can be factorized using a **QR factorization**, with column interchanges, as

$$A = Q \begin{pmatrix} R \\ 0 \end{pmatrix} P^T,$$

where Q is an m by m orthogonal matrix, R is an r by n upper trapezoidal matrix and P is an n by n permutation matrix. The pseudo-inverse of A is then given by

$$A^+ = PR^T(RR^T)^{-1}\tilde{Q}^T,$$

where \tilde{Q} consists of the first r columns of Q .

- (d) if $\text{rank}(A) = r \leq k = \min(m, n)$, then A can be factorized as the **singular value decomposition**

$$A = U\Sigma V^T,$$

where U is an m by m orthogonal matrix, V is an n by n orthogonal matrix and Σ is an m by n diagonal matrix with non-negative diagonal elements σ . The first k columns of U and V are the **left-** and **right-hand singular vectors** of A respectively and the k diagonal elements of Σ are the **singular values** of A . Σ may be chosen so that

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_k \geq 0$$

and in this case if $\text{rank}(A) = r$ then

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0, \quad \sigma_{r+1} = \cdots = \sigma_k = 0.$$

If \tilde{U} and \tilde{V} consist of the first r columns of U and V respectively and $\tilde{\Sigma}$ is an r by r diagonal matrix with diagonal elements $\sigma_1, \sigma_2, \dots, \sigma_r$ then A is given by

$$A = \tilde{U}\tilde{\Sigma}\tilde{V}^T$$

and the pseudo-inverse of A is given by

$$A^+ = \tilde{V}\tilde{\Sigma}^{-1}\tilde{U}^T.$$

Notice that

$$A^T A = V(\Sigma^T \Sigma)V^T$$

which is the classical eigenvalue (spectral) factorization of $A^T A$.

- (e) if A is complex then the above relationships are still true if we use ‘unitary’ in place of ‘orthogonal’ and conjugate transpose in place of transpose. For example, the singular value decomposition of A is

$$A = U\Sigma V^H,$$

where U and V are unitary, V^H the conjugate transpose of V and Σ is as in (d) above.

2.2 Matrix Factorizations

The routines in this section perform matrix factorizations which are required for the solution of systems of linear equations with various special structures. A few routines which perform associated computations are also included.

Other routines for matrix factorizations are to be found in Chapters F07, F08 and F11.

This section also contains a few routines associated with eigenvalue problems (see Chapter F02). (Historical note: this section used to contain many more such routines, but they have now been superseded by routines in Chapter F08.)

2.3 Matrix Arithmetic and Manipulation

The intention of routines in this section (sub-chapters F01C, F01V and F01Z) is to cater for some of the commonly occurring operations in matrix manipulation, e.g., transposing a matrix or adding part of one matrix to another, and for conversion between different storage formats, e.g., conversion between rectangular band matrix storage and packed band matrix storage. For vector or matrix-vector or matrix-matrix operations refer to Chapters F06 and F16.

2.4 Matrix Functions

Given a square matrix A , the matrix function $f(A)$ is a matrix with the same dimensions as A which provides a generalization of the scalar function f .

If A has a full set of eigenvectors V then A can be factorized as

$$A = VDV^{-1},$$

where D is the diagonal matrix whose diagonal elements, d_i , are the eigenvalues of A . $f(A)$ is given by

$$f(A) = Vf(D)V^{-1},$$

where $f(D)$ is the diagonal matrix whose i th diagonal element is $f(d_i)$.

In general, A may not have a full set of eigenvectors. The matrix function can then be defined via a Cauchy integral. For $A \in \mathbb{C}^{n \times n}$,

$$f(A) = \frac{1}{2\pi i} \int_{\Gamma} f(z)(zI - A)^{-1} dz,$$

where Γ is a closed contour surrounding the eigenvalues of A , and f is analytic within Γ .

Algorithms for computing matrix functions are usually tailored to a specific function. Currently Chapter F01 contains routines for calculating the exponential, logarithm, sine, cosine, sinh and cosh of both real and complex matrices. In addition there are routines to compute a general function of real symmetric and complex Hermitian matrices and a general function of general real and complex matrices.

The condition number of a matrix function is a measure of its sensitivity to perturbations in the data. Chapter F01 contains routines for estimating the condition number of the matrix exponential, logarithm, sine, cosine, sinh or cosh for real or complex matrices. It also contains routines for estimating the condition number of a general function of a real or complex matrix.

3 Recommendations on Choice and Use of Available Routines

3.1 Matrix Inversion

Note: before using any routine for matrix inversion, consider carefully whether it is really needed.

Although the solution of a set of linear equations $Ax = b$ can be written as $x = A^{-1}b$, the solution should **never** be computed by first inverting A and then computing $A^{-1}b$; the routines in Chapters F04 or F07 should **always** be used to solve such sets of equations directly; they are faster in execution, and numerically more stable and accurate. Similar remarks apply to the solution of least squares problems which again should be solved by using the routines in Chapters F04 and F08 rather than by computing a pseudo-inverse.

(a) Nonsingular square matrices of order n

This chapter describes techniques for inverting a general real matrix A and matrices which are positive definite (have all eigenvalues positive) and are either real and symmetric or complex and Hermitian. It is wasteful and uneconomical not to use the appropriate routine when a matrix is known to have one of these special forms. A general routine must be used when the matrix is not known to be positive definite. In most routines the inverse is computed by solving the linear equations $Ax_i = e_i$, for $i = 1, 2, \dots, n$, where e_i is the i th column of the identity matrix.

Routines are given for calculating the approximate inverse, that is solving the linear equations just once, and also for obtaining the accurate inverse by successive iterative corrections of this first approximation. The latter, of course, are more costly in terms of time and storage, since each correction involves the solution of n sets of linear equations and since the original A and its LU decomposition must be stored together with the first and successively corrected approximations to the inverse. In practice the storage requirements for the ‘corrected’ inverse routines are about double those of the ‘approximate’ inverse routines, though the extra computer time is not prohibitive since the same matrix and the same LU decomposition is used in every linear equation solution.

Despite the extra work of the ‘corrected’ inverse routines they are superior to the ‘approximate’ inverse routines. A correction provides a means of estimating the number of accurate figures in the

inverse or the number of ‘meaningful’ figures relating to the degree of uncertainty in the coefficients of the matrix.

The residual matrix $R = AX - I$, where X is a computed inverse of A , conveys useful information. Firstly $\|R\|$ is a bound on the relative error in X and secondly $\|R\| < \frac{1}{2}$ guarantees the convergence of the iterative process in the ‘corrected’ inverse routines.

The decision trees for inversion show which routines in Chapter F04 and Chapter F07 should be used for the inversion of other special types of matrices not treated in the chapter.

(b) General real rectangular matrices

For real matrices F08AEF (DGEQRF) and F01QJF return QR and RQ factorizations of A respectively and F08BFF (DGEQP3) returns the QR factorization with column interchanges. The corresponding complex routines are F08ASF (ZGEQRF), F01RJF and F08BTF (ZGEQP3) respectively. Routines are also provided to form the orthogonal matrices and transform by the orthogonal matrices following the use of the above routines. F01QGF and F01RGF form the RQ factorization of an upper trapezoidal matrix for the real and complex cases respectively.

F01BLF uses the QR factorization as described in Section 2.1(ii)(a) and is the only routine that explicitly returns a pseudo-inverse. If $m \geq n$, then the routine will calculate the pseudo-inverse A^+ of the matrix A . If $m < n$, then the n by m matrix A^T should be used. The routine will calculate the pseudo-inverse $Z = (A^T)^+ = (A^+)^T$ of A^T and the required pseudo-inverse will be Z^T . The routine also attempts to calculate the rank, r , of the matrix given a tolerance to decide when elements can be regarded as zero. However, should this routine fail due to an incorrect determination of the rank, the singular value decomposition method (described below) should be used.

F08KBF (DGESVD) and F08KPF (ZGESVD) compute the singular value decomposition as described in Section 2.2 for real and complex matrices respectively. If A has rank $r \leq k = \min(m, n)$ then the $k - r$ smallest singular values will be negligible and the pseudo-inverse of A can be obtained as $A^+ = V\Sigma^{-1}U^T$ as described in Section 2. If the rank of A is not known in advance it can be estimated from the singular values (see Section 2.4 in the F04 Chapter Introduction). In the real case with $m \geq n$, F02WDF provides details of the QR factorization or the singular value decomposition depending on whether or not A is of full rank and for some problems provides an attractive alternative to F08KBF (DGESVD). For large sparse matrices, leading terms in the singular value decomposition can be computed using routines from Chapter F12.

3.2 Matrix Factorizations

Each of these routines serves a special purpose required for the solution of sets of simultaneous linear equations or the eigenvalue problem. For further details you should consult Sections 3 or 4 in the F02 Chapter Introduction or Sections 3 or 4 in the F04 Chapter Introduction.

F01BRF and F01BSF are provided for factorizing general real sparse matrices. A more recent algorithm for the same problem is available through F11MEF. For factorizing real symmetric positive definite sparse matrices, see F11JAF. These routines should be used only when A is **not** banded and when the total number of nonzero elements is less than 10% of the total number of elements. In all other cases either the band routines or the general routines should be used.

3.3 Matrix Arithmetic and Manipulation

The routines in the F01C section are designed for the general handling of m by n matrices. Emphasis has been placed on flexibility in the parameter specifications and on avoiding, where possible, the use of internally declared arrays. They are therefore suited for use with large matrices of variable row and column dimensions. Routines are included for the addition and subtraction of sub-matrices of larger matrices, as well as the standard manipulations of full matrices. Those routines involving matrix multiplication may use additional-precision arithmetic for the accumulation of inner products. See also Chapter F06.

The routines in the F01V (LAPACK) and F01Z section are designed to allow conversion between full storage format and one of the packed storage schemes required by some of the routines in Chapters F02, F04, F06, F07 and F08.

3.3.1 NAG Names and LAPACK Names

Routines with NAG name beginning F01V may be called either by their NAG names or by their LAPACK names. When using the NAG Library, the double precision form of the LAPACK name must be used (beginning with D- or Z-).

References to Chapter F01 routines in the manual normally include the LAPACK double precision names, for example, F01VEF (DTRTTF).

The LAPACK routine names follow a simple scheme (which is similar to that used for the BLAS in Chapter F06). Most names have the structure XYYTZZ, where the components have the following meanings:

- the initial letter, X, indicates the data type (real or complex) and precision:
 - S – real, single precision (in Fortran, 4 byte length REAL)
 - D – real, double precision (in Fortran, 8 byte length REAL)
 - C – complex, single precision (in Fortran, 8 byte length COMPLEX)
 - Z – complex, double precision (in Fortran, 16 byte length COMPLEX)
- the fourth letter, T, indicates that the routine is performing a storage scheme transformation (conversion)
- the letters YY indicate the original storage scheme used to store a triangular part of the matrix A , while the letters ZZ indicate the target storage scheme of the conversion (YY cannot equal ZZ since this would do nothing):
 - TF – Rectangular Full Packed Format (RFP)
 - TP – Packed Format
 - TR – Full Format

3.4 Matrix Functions

F01ECF and F01FCF compute the matrix exponential, e^A , of a real and complex square matrix A respectively. If estimates of the condition number of the matrix exponential are required then F01JAF and F01KAF should be used.

F01EDF and F01FDF compute the matrix exponential, e^A , of a real symmetric and complex Hermitian matrix respectively. If the matrix is real symmetric, or complex Hermitian then it is recommended that F01EDF, or F01FDF be used as they are more efficient and, in general, more accurate than F01ECF and F01FCF.

F01EJF and F01FJF compute the principal matrix logarithm, $\log(A)$, of a real and complex square matrix A respectively. If estimates of the condition number of the matrix logarithm are required then F01JAF and F01KAF should be used.

F01EKF and F01FKF compute the matrix exponential, sine, cosine, sinh or cosh of a real and complex square matrix A respectively. If the matrix exponential is required then it is recommended that F01ECF or F01FCF be used as they are, in general, more accurate than F01EKF and F01FKF. If estimates of the condition number of the matrix function are required then F01JAF and F01KAF should be used.

F01ELF and F01EMF compute the matrix function, $f(A)$, of a real square matrix. F01FLF and F01FMF compute the matrix function of a complex square matrix. The derivatives of f are required for these computations. F01ELF and F01FLF use numerical differentiation to obtain the derivatives of f . F01EMF and F01FMF use derivatives you have supplied. If estimates of the condition number of the matrix function are required and you are supplying derivatives of f , then F01JCF and F01KCF should be used. If estimates of the condition number are required but you are not supplying derivatives then F01JBF and F01KBF should be used.

F01EFF and F01FFF compute the matrix function, $f(A)$, of a real symmetric and complex Hermitian matrix A respectively. If the matrix is real symmetric or complex Hermitian then it is recommended that F01EFF or F01FFF be used as they are more efficient and, in general, more accurate than F01ELF, F01EMF, F01FLF and F01FMF.


F01GAF and F01HAF compute the matrix function $e^{tA}B$ for explicitly stored dense real and complex matrices A and B respectively while F01GBF and F01HBF compute the same using reverse communication. In the latter case, control is returned to you. You should calculate any required matrix-matrix products and then call the routine again.

4 Decision Trees

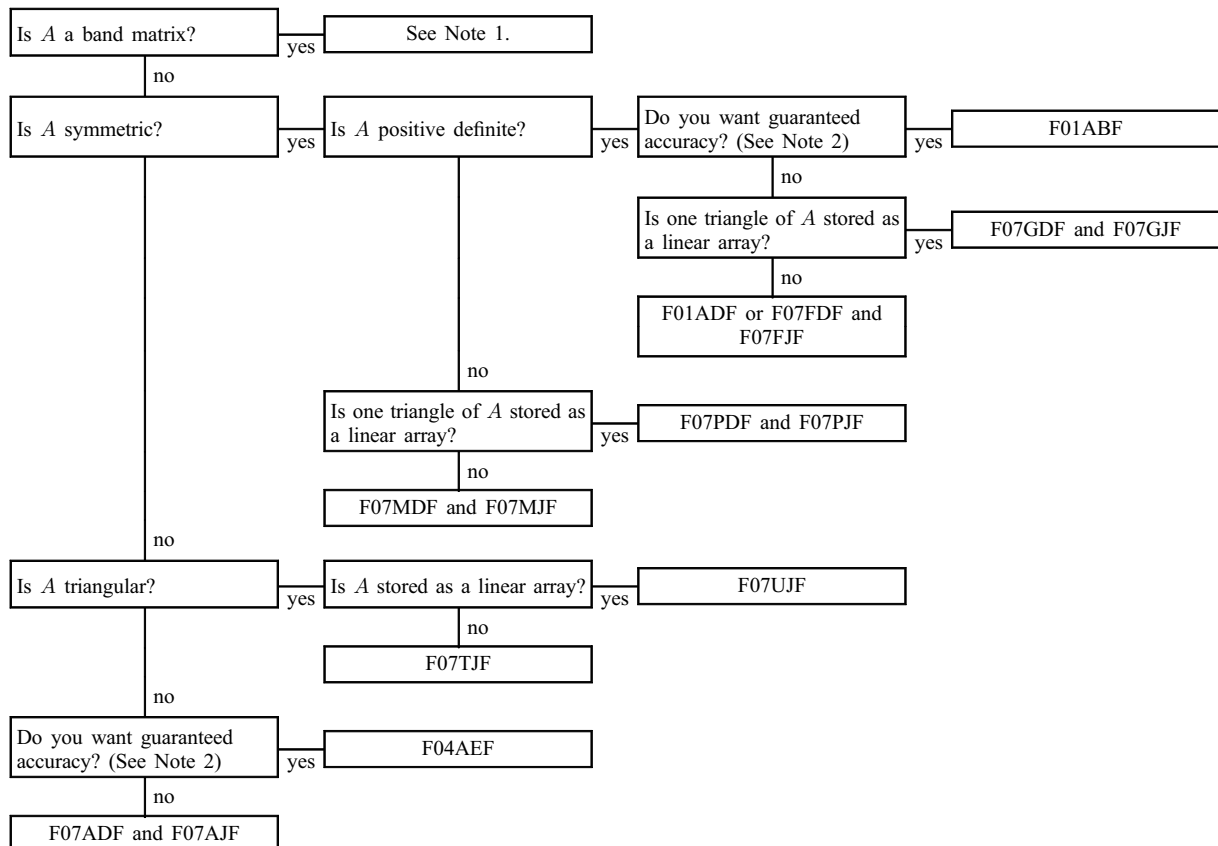
The decision trees show the routines in this chapter and in Chapter F04, Chapter F07 and Chapter F08 that should be used for inverting matrices of various types. They also show which routine should be used to calculate various matrix functions.

(i) Matrix Inversion:

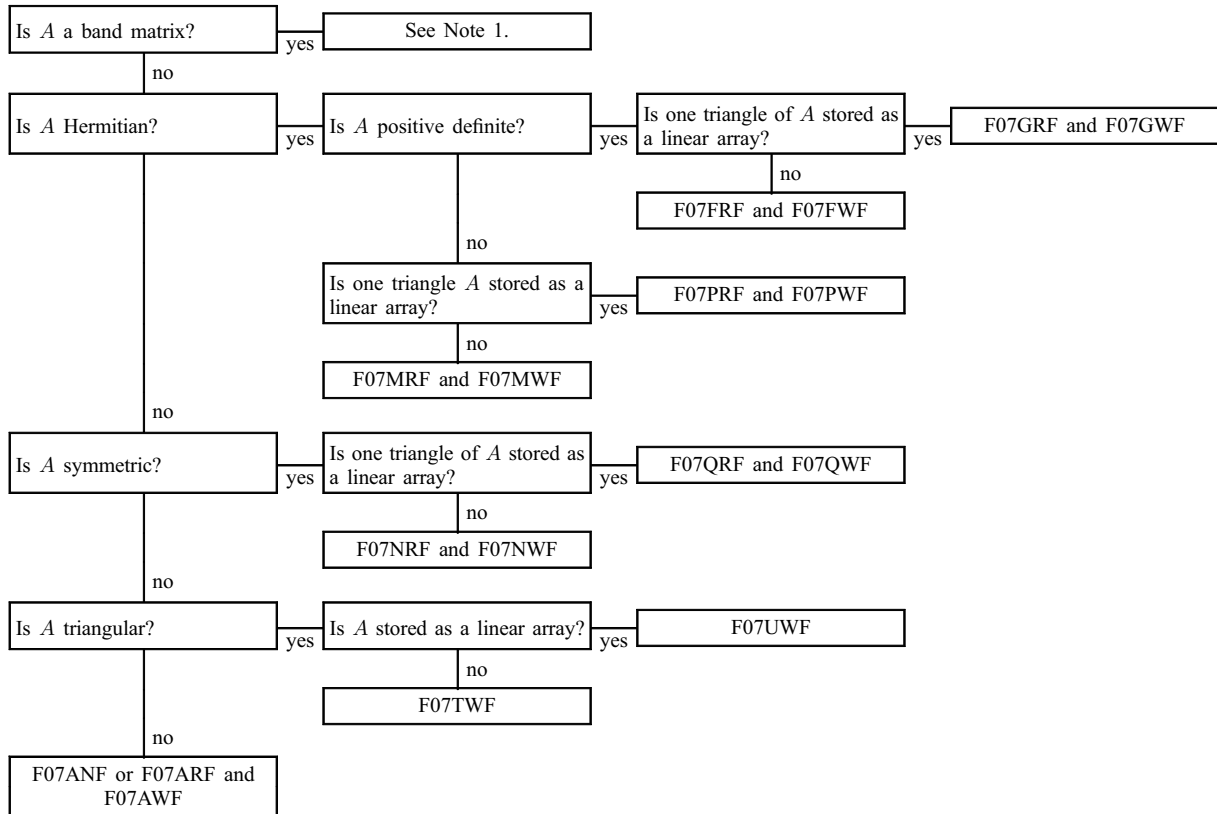
Tree 1



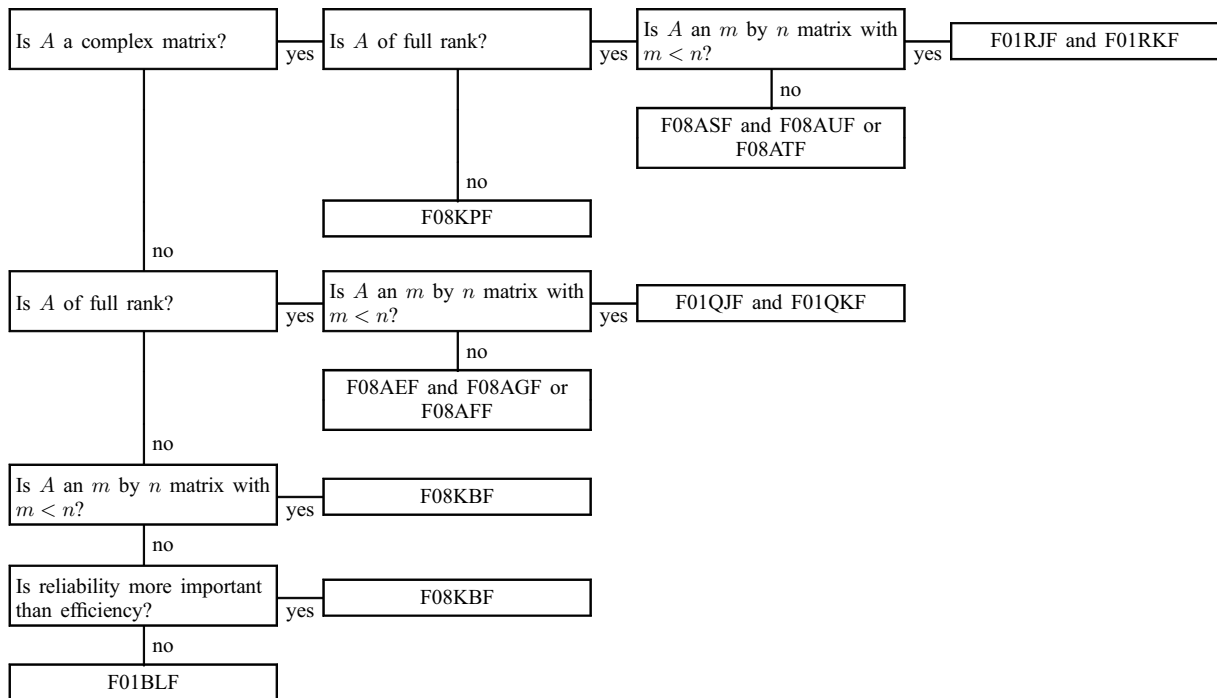
Tree 2: Inverse of a real n by n matrix of full rank



Tree 3: Inverse of a complex n by n matrix of full rank



Tree 4: Pseudo-inverses



Note 1: the inverse of a band matrix A does not in general have the same shape as A , and no routines are provided specifically for finding such an inverse. The matrix must either be treated as a full matrix, or the equations $AX = B$ must be solved, where B has been initialized to the identity matrix I . In the latter case, see the decision trees in Section 4 in the F04 Chapter Introduction.

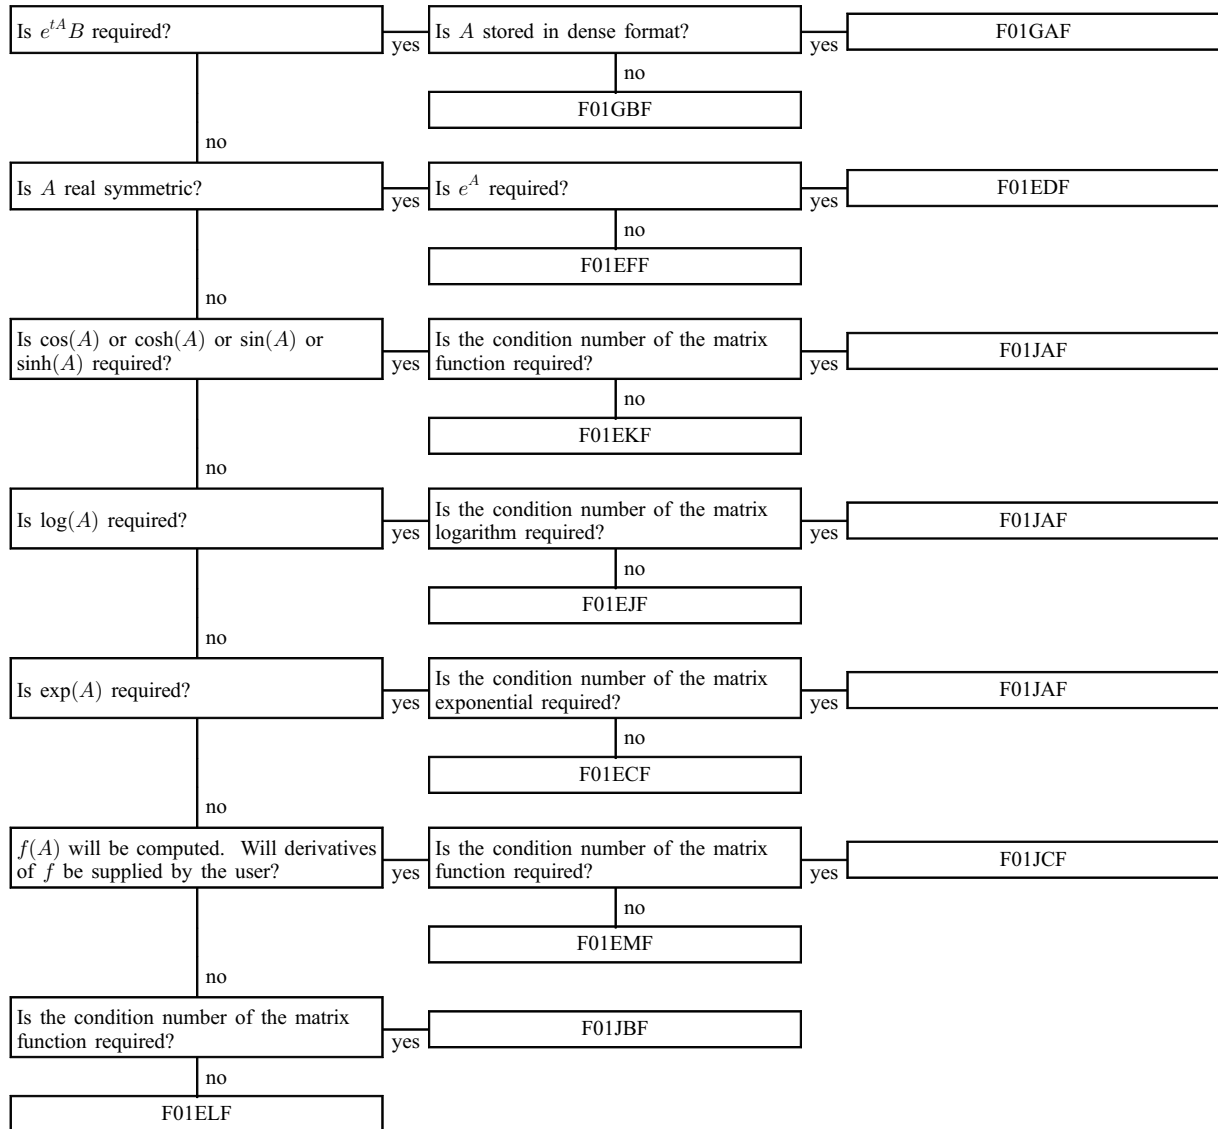
Note 2: by ‘guaranteed accuracy’ we mean that the accuracy of the inverse is improved by use of the iterative refinement technique using additional precision.

(ii) **Matrix Factorizations:** see the decision trees in Section 4 in the F02 and F04 Chapter Introductions.

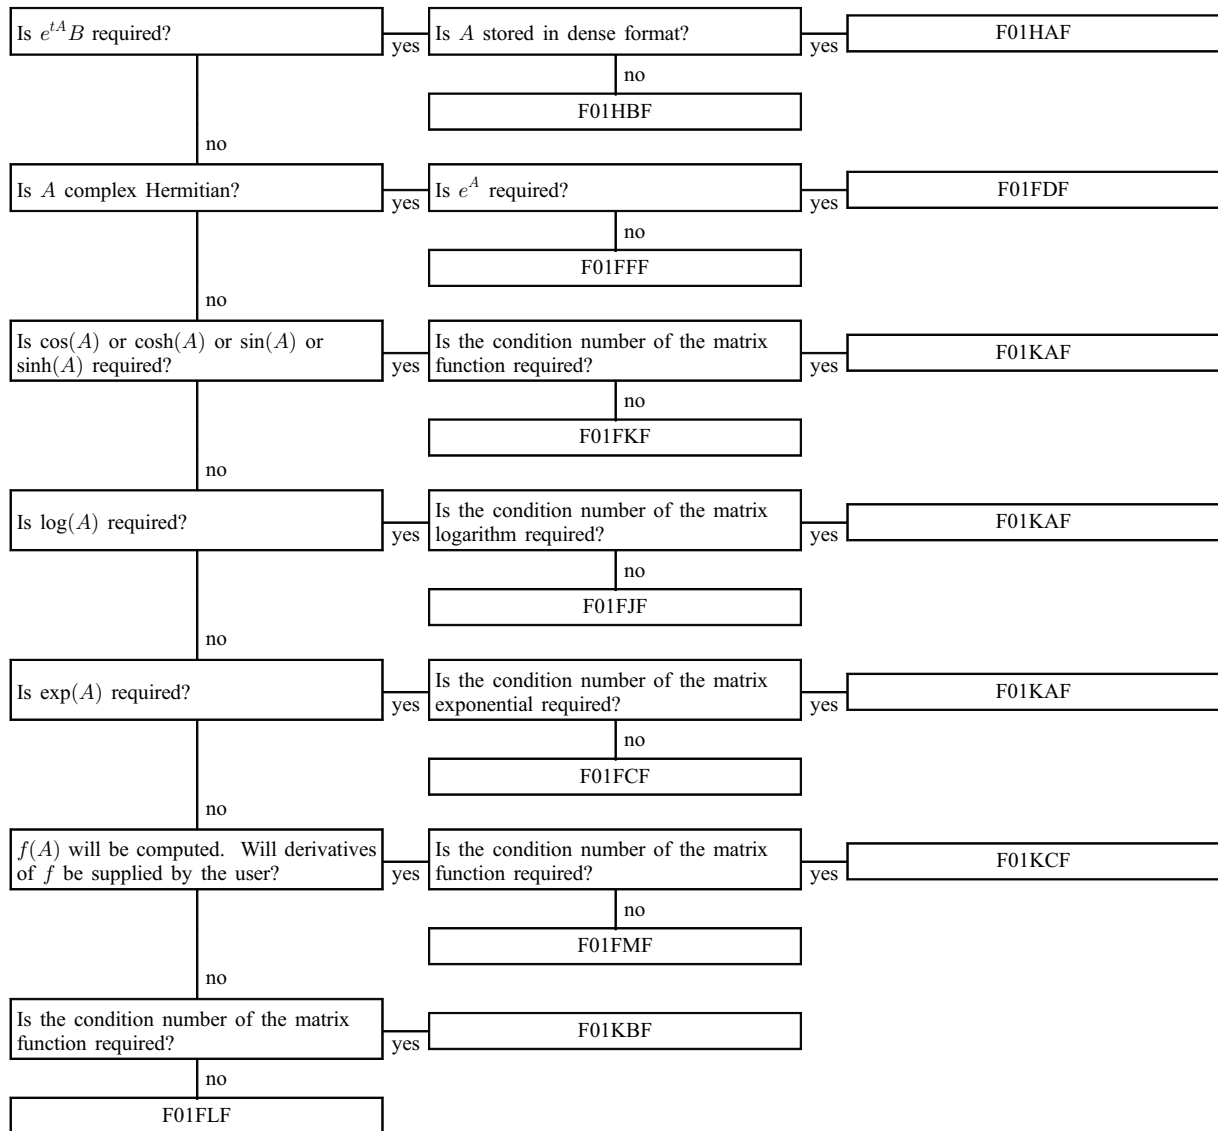
(iii) **Matrix Arithmetic and Manipulation:** not appropriate.

(iv) **Matrix Functions:**

Tree 5: Matrix functions $f(A)$ of an n by n real matrix A



Tree 6: Matrix functions $f(A)$ of an n by n complex matrix A



5 Functionality Index

Action of the matrix exponential on a complex matrix	F01HAF
Action of the matrix exponential on a complex matrix (reverse communication) ..	F01HBF
Action of the matrix exponential on a real matrix	F01GAF
Action of the matrix exponential on a real matrix (reverse communication)	F01GBF
Inversion (also see Chapter F07), real m by n matrix, pseudo-inverse.....	F01BLF
real symmetric positive definite matrix, accurate inverse	F01ABF
approximate inverse.....	F01ADF

Matrix Arithmetic and Manipulation,	
matrix addition,	
complex matrices	F01CWF
real matrices	F01CTF
matrix multiplication	F01CKF
matrix storage conversion,	
full to packed triangular storage,	
complex matrices	F01VBF (ZTRTTP)
real matrices	F01VAF (DTRTTP)
full to Rectangular Full Packed storage,	
complex matrix	F01VFF (ZTRTTF)
real matrix	F01VEF (DTRTTF)
packed band ↔ rectangular storage, special provision for diagonal	
complex matrices	F01ZDF
real matrices	F01ZCF
packed triangular to full storage,	
complex matrices	F01VDF (ZTPTTR)
real matrices	F01VCF (DTPTTR)
packed triangular to Rectangular Full Packed storage,	
complex matrices	F01VKF (ZTPTTF)
real matrices	F01VJF (DTPTTF)
packed triangular ↔ square storage, special provision for diagonal	
complex matrices	F01ZBF
real matrices	F01ZAF
Rectangular Full Packed to full storage,	
complex matrices	F01VHF (ZTFTR)
real matrices	F01VGF (DTFTR)
Rectangular Full Packed to packed triangular storage,	
complex matrices	F01VMF (ZFTTTP)
real matrices	F01VLF (DTFTTP)
matrix subtraction,	
complex matrices	F01CWF
real matrices	F01CTF
matrix transpose	F01CRF
Matrix function,	
complex Hermitian n by n matrix,	
matrix exponential	F01FDF
matrix function	F01FFF
complex n by n matrix,	
condition number for a matrix exponential, logarithm, sine, cosine, sinh or cosh	F01KAF
condition number for a matrix function, using numerical differentiation	F01KBF
condition number for a matrix function, using user-supplied derivatives	F01KCF
matrix exponential	F01FCF
matrix exponential, sine, cosine, sinh or cosh	F01FKF
matrix function, using numerical differentiation	F01FLF
matrix function, using user-supplied derivatives	F01FMF
matrix logarithm	F01FJF
real n by n matrix,	
condition number for a matrix function, using numerical differentiation	F01JBF
condition number for a matrix function, using user-supplied derivatives	F01JCF
condition number for the matrix exponential, logarithm, sine, cosine, sinh or cosh	F01JAF
matrix exponential	F01ECF
matrix exponential, sine, cosine, sinh or cosh	F01EKF
matrix function, using numerical differentiation	F01ELF
matrix function, using user-supplied derivatives	F01EMF

matrix logarithm	F01EJF
real symmetric n by n matrix, matrix exponential	F01EDF
matrix function	F01EFF
Matrix Transformations,	
complex matrix, form unitary matrix	F01RKF
complex m by n ($m \leq n$) matrix, RQ factorization	F01RJF
complex upper trapezoidal matrix, RQ factorization	F01RGF
eigenproblem $Ax = \lambda Bx$, A , B banded, reduction to standard symmetric problem	F01BVF
real almost block-diagonal matrix, LU factorization	F01LHF
real band symmetric positive definite matrix, $ULDL^T U^T$ factorization	F01BUF
variable bandwidth, LDL^T factorization	F01MCF
real matrix, form orthogonal matrix	F01QKF
real m by n ($m \leq n$) matrix, RQ factorization	F01QJF
real sparse matrix, factorization	F01BRF
factorization, known sparsity pattern	F01BSF
real upper trapezoidal matrix, RQ factorization	F01QGF
tridiagonal matrix, LU factorization	F01LEF

6 Auxiliary Routines Associated with Library Routine Parameters

None.

7 Routines Withdrawn or Scheduled for Withdrawal

The following lists all those routines that have been withdrawn since Mark 17 of the Library or are scheduled for withdrawal at one of the next two marks.

Withdrawn Routine	Mark of Withdrawal	Replacement Routine(s)
F01AAF	17	F07ADF (DGETRF) and F07AJF (DGETRI)
F01AEF	18	F07FDF (DPOTRF), F08SEF (DSYGST) and F06EGF (DSWAP)
F01AFF	18	F06EGF (DSWAP) and F06YJF (DTRSM)
F01AGF	18	F08FEF (DSYTRD)
F01AHF	18	F08FGF (DORMTR)
F01AJF	18	F08FEF (DSYTRD) and F08FFF (DORGTR)
F01AKF	18	F08NEF (DGEHRD)
F01ALF	18	F08NGF (DORMHR)
F01AMF	18	F08NSF (ZGEHRD)
F01ANF	18	F08NUF (ZUNMHR)
F01APF	18	F06QFF and F08NFF (DORGHR)
F01ATF	18	F08NHF (DGEBAL)
F01AUF	18	F08NHF (DGEBAL)
F01AVF	18	F08NVF (ZGEBAL)
F01AWF	18	F08NWF (ZGEBAL)
F01AXF	18	F06EFF (DCOPY) and F08BEF (DGEQPF)

F01AYF	18	F08GEF (DSPTRD)
F01AZF	18	F08GGF (DOPMTR)
F01BCF	18	F08FSF (ZHETRD) and F08FTF (ZUNGTR)
F01BDF	18	F07FDF (DPOTRF), F08SEF (DSYGST) and F06EGF (DSWAP)
F01BEF	18	F06YFF (DTRMM) and F06EGF (DSWAP)
F01BNF	17	F07FRF (ZPOTRF)
F01BPF	17	F07FRF (ZPOTRF) and F07FWF (ZPOTRI)
F01BTF	18	F07ADF (DGETRF)
F01BWF	18	F08HEF (DSBTRD)
F01BXF	17	F07FDF (DPOTRF)
F01LBF	18	F07BDF (DGBTRF)
F01MAF	19	F11JAF
F01NAF	17	F07BRF (ZGBTRF)
F01QCF	18	F08AEF (DGEQRF)
F01QDF	18	F08AGF (DORMQR)
F01QEF	18	F08AFF (DORGQR)
F01QFF	18	F08BEF (DGEQPF)
F01RCF	18	F08ASF (ZGEQRF)
F01RDF	18	F08AUF (ZUNMQR)
F01REF	18	F08ATF (ZUNGQR)
F01RFF	18	F08BSF (ZGEQPF)

8 References

Golub G H and Van Loan C F (1996) *Matrix Computations* (3rd Edition) Johns Hopkins University Press, Baltimore

Higham N J (2008) *Functions of Matrices: Theory and Computation* SIAM, Philadelphia, PA, USA

Wilkinson J H (1965) *The Algebraic Eigenvalue Problem* Oxford University Press, Oxford

Wilkinson J H (1977) Some recent advances in numerical linear algebra *The State of the Art in Numerical Analysis* (ed D A H Jacobs) Academic Press

Wilkinson J H and Reinsch C (1971) *Handbook for Automatic Computation II, Linear Algebra* Springer-Verlag