

JDQR

`jdqr` computes eigenpairs of a square matrix or operator.

`Lambda = jdqr(A)` returns the absolute largest eigenvalues of `A` in a `k` vector `Lambda`. Here `k = min(5,n)` and `n = size(A,1)`. `jdqr(A)` (without output argument) displays the `k` eigenvalues.

`[X,Jordan] = jdqr(A,B)` returns the eigenvectors `X` and the Jordan structure `Jordan`: `A*X = X*Jordan`. The diagonal of `Jordan` contains the eigenvalues: `Lambda = diag(Jordan)`. `Jordan` is an `k` by `k` matrix with the eigenvalues on the diagonal and possibly non-zeros on the first upper diagonal elements. The other entries are zero. The columns of `X` have norm 1.

`[X,Jordan,Q,S] = jdqr(A)`

If four or more output arguments are required then `Q` is `n` by `k` orthonormal, and `S` is `k` by `k` upper triangular such that they form a partial generalized Schur decomposition: `A*Q = Q*S`. Then `Lambda = diag(S)` and `X = Q*Y` with `Y` the eigenvectors of the pair `S`: `S*Y = Y*Jordan` (see also `Options.Schur`).

```
... = jdqr(A)
... = jdqr('Afun')
... = jdqr('Afun',n)
```

The first input argument is either a square matrix (which can be full or sparse, symmetric or nonsymmetric, real or complex), or a string containing the name of an M-file which applies a linear operator to the columns of a given matrix. In the latter case, the M-file, say `Afun.m`, must return the dimension `n` of the problem with `n = Afun([],'dimension')` or `n` must be specified in the list of input arguments. For example, `jdqr('fft',...)` is much faster than `jdqr(F,...)`, where `F` is the explicit FFT matrix.

The remaining input arguments are optional and can be given in practically any order:

```
... = jdqr(A,k,Sigma,Options,M)
... = jdqr('Afun',k,Sigma,Options,M),
```

where

<code>k</code>	an integer, the number of desired eigenvalues.
<code>Sigma</code>	a scalar shift or a two letter string.
<code>Options</code>	a structure containing additional parameters.
<code>M</code>	a string or a matrix that specifies the preconditioner.

If `k` is not specified, then `k = min(n,5)` eigenvalues are computed.

If `Sigma` is not specified, then the `k`th eigenvalues largest in magnitude are computed. If `Sigma` is a real or complex scalar, then the `k`th eigenvalues nearest `Sigma` are computed. If `Sigma` is row vector of size `(1,m)`, then the `j`th eigenvalue nearest to `Sigma(1,min(m,j))` is computed for `j = 1:k`. `Sigma` is the “target” for the desired eigenvalues. If `Sigma` is one of the following strings, then it specifies the desired eigenvalues.

Sigma Specified eigenvalues
 'LM' Largest Magnitude
 'SM' Smallest Magnitude (same as **Sigma** = 0)
 'LR' Largest Real part
 'SR' Smallest Real part
 'BE' Both Ends. Computes $k/2$ eigenvalues from each end of the spectrum (one more from the high end if k is odd.)

If 'TestSpace' is 'Harmonic' (see Options), then **Sigma** = 0 is the default, otherwise **Sigma** = 'LM' is the default.

The Options structure specifies certain parameters in the algorithm.

Field name	Parameter	Default
Options.Tol	Convergence tolerance: $\text{norm}(r) \leq \text{tol}/\sqrt{k}$	1e-8
Options.jmin	Minimum dimension search subspace V	k+5
Options.jmax	Maximum dimension search subspace V	jmin+5
Options.MaxIt	Maximum number of iterations.	100
Options.v0	Starting space	ones+0.1*rand
Options.Schur	Gives schur decomposition If 'yes', then X and Jordan are not computed and [Q,S,history] is the list of output arguments.	'no'
Options.TestSpace	Defines the test subspace W 'Standard': $W = V$ 'Harmonic': $W = A*V - \text{sigma}*V$	'standard'
Options.Disp	If Disp=1 or Disp='yes', then, the input parameters, the residual size at each step, and the eigenvalues at detection are displayed, and the convergence history is plotted. If Disp=2 then, in addition, approximate eigenvalues are plotted at each step.	'no'
Options.LSsolver	Linear solver	'GMRES'
Options.LS_Tol	Residual reduction linear solver	1,0.7,0.7 ² ,...
Options.LS_MaxIt	Maximum number it. linear solver	5
Options.LS_ell	ell for BiCGstab(ell)	4
Options.Precond	Preconditioner(see below).	identity.

For instance,

```
Options = struct('Tol',1.0e-8,'LSsolver','BiCGstab','LS_ell',4,'Precond',M);
```

changes the convergence tolerance to 1.0e-8, takes BiCGstab as linear solver, and takes M as preconditioner (for ways of defining M, see below).

There are a few other Options that can be specified. They are listed below .

```
[X,Jordan,history] = jdqr(A,...)
```

```
[X,Jordan,Q,S,history] = jdqr(A,...)
```

returns also the convergence history.

history is an array of 3 columns: history(i,1) is the residual norm at step $j = \text{history}(i,2)$, history(i,2), history(i,3) is the cumulative number of multiplications by A at step j . If a search for a new eigenvalue is started at step J then $j = \text{history}(i,2) = \text{history}(i+1,2)$, history(i,1) is the norm of the "old" residual, and history(i+1,1) is the norm of the "new" one. history is empty if the required number of eigenvalues are detected in the initialization phase.

jdqr without input arguments returns the options and its defaults.

Preconditioning in jdqr

The action ‘M inverse’ of the preconditioner M (an approximation of $A - \lambda I$) on an n-vector v can be defined in the Options

```
Options.Precond
Options.L_Precond      (same as Options.Precond)
Options.U_Precond
Options.P_Precond
```

and also in the argument list:

```
... = jdqr(...,k,sigma,M,options)
... = jdqr(...,k,sigma,L,U,options)
... = jdqr(...,k,sigma,'M',options)
... = jdqr(...,k,sigma,'L','U',options)
```

If no preconditioner has been specified (or is []), then $M \setminus v = v$ (M is the identity).

If Precond is an n by n matrix, say, K, then

$$M \setminus v = K \setminus v.$$

If Precond is an N by 2*N matrix, say, K, then

$$M \setminus v = U \setminus L \setminus v, \text{ where } K = [L, U], \text{ and } L \text{ and } U \text{ are } n \text{ by } n \text{ matrices.}$$

If Precond is a string, say, 'Mi', then

if $Mi(v, 'L')$ and $Mi(v, 'U')$ return n-vectors

$$M \setminus v = Mi(Mi(v, 'L'), 'U')$$

otherwise

$$M \setminus v = Mi(v) \text{ or } M \setminus v = Mi(v, 'preconditioner').$$

Note that Precond and A can be the same string.

If L_Precond and U_Precond are strings, say, 'Li' and 'Ui', respectively, then

$$M \setminus v = Ui(Li(v)).$$

If (P_precond,) L_Precond, and U_precond are n by n matrices, say, (P,) L, and U, respectively, then

$$M \setminus v = U \setminus L \setminus (P * v) \quad (\text{i.e., } P * M = L * U).$$

The way the preconditioner is used can be specified in the Options, see below.

Way of using a preconditioner

The way the preconditioner is used can be specified in the Options.

```
Options.Type_Precond (default 'left')
```

The preconditioner can be used as explicit left preconditioner ('left'), as explicit right preconditioner ('right'), or implicitly ('impl').

In this subsection,

- an MV (matrix vector multiplication) is an operation by A,
- a PS (preconditioner solve) is a solution of the system $Mt = v$, where M is the preconditioner (i.e., the computation of $M \setminus v$),

Explicit versus implicit. If explicit preconditioning is used, then there is an additional PS needed each time the correction equation is solved. The total number of PSs that `jdqr` will take is equal to the total number of MVs plus the number of Jacobi-Davidson steps (the number of outer iterations).

With implicit preconditioning the number of PSs reduces to the number of MVs plus the number of detected eigenvalues. However, implicit preconditioning requires more memory. For `BiCGstab(ell)`, $2 \cdot \text{ell}$ additional `n`-vectors have to be stored. If GMRES is requested as linear solver, then FGMRES is used, requiring storage of an additional m `n`-vectors. Here m is the maximum number of steps that GMRES needs to achieve the required residual reduction.

Implicit preconditioning in CG does not lead to additional memory requirements and `jdqr` uses implicit preconditioning whenever CG is selected as linear solver.

Implicit preconditioning in MINRES and in SYMMLQ does not reduce the PSs. Moreover, it requires storage of one additional `n`-vector. However, with implicit preconditioning, there is no need specify the factors `L` and `L'` of the preconditioner `M` ($M = LL^*$). Therefore, `jdqr` uses implicit preconditioning also if MINRES or SYMMLQ is selected as linear solver.

The methods CG, MINRES and SYMMLQ require a positive definite preconditioner. By storing the preconditioned vectors of the search subspace the number of PSs can be reduced even further. Then the number of PSs will be equal to the number of MVs. However, this strategy has not been implemented in `jdqr`.

There may be a slight deviation in the count of PSs and MVs if `jdqr` detects more than one eigenpair at the same iteration step.

Right versus left. If explicit right preconditioning is used, then the size of the residual is available in the inner loop, that is, the size of the residual of the iterative solver for the correction equation. With explicit left preconditioning, only 'preconditioned' residuals are available. However, right preconditioning, as well as implicit preconditioning, requires slightly more projections in GMRES and `BiCGstab(ell)`.

Additional Options

`Options.Pairs` (default 'no')

If 'yes', then `jdqr` searches for the complex conjugate eigenpair whenever an eigenpair has been detected. If `A` and `B` are real matrices, or the operators correspond to real matrices, then $(\bar{\lambda}, \bar{x})$ is an eigenpair if (λ, x) is one. Since the eigenpairs are not computed in full accuracy and since a generalized Schur decomposition is computed instead of eigenpairs, the conjugate of an approximate eigenpair may not have the required precision and `jdqr` may take additional iterations to obtain the conjugate pair in the desired accuracy.

`Options.FixShift` (default 'no')

If `Options.FixShift` is scalar and `Sigma(1,:)` is a scalar, then `jdqr` takes `Sigma` as shift in the correction equation until the norm of the residual times `Options.FixShift` is less than 1. From then on, the shift is taken equal to the present approximate eigenvalue.

If `Options.FixShift = 'yes'` then `Options.FixShift = 1.0e+3`.

If `Options.FixShift = 'no'` is the same as `Options.FixShift = 0`.

`Options.Track` (default '1e-4')

If the wanted eigenvalue is relatively far from the target, then the algorithm may select approximate eigenvalues that are accidentally close to the target instead of the approximate eigenvalue that is close to the wanted eigenvalue. To avoid this type of misselection, the target can be moved to an approximate eigenvalue that is close to the wanted eigenvalue. The size of the norm of the residual is used to measure the quality of the approximate eigenvalue. If $\text{norm}(\mathbf{r}) \leq \text{Options.track}$, then the associated approximate eigenvalue is used as target in the next iteration step.

`Options.AvoidStag` (default 'no')

In some situations, the algorithm stagnates because the computed expansion vector for the search subspace belongs to the search subspace or is close to it. With 'yes', `jdqr` tries to remedy this type of stagnation. In the correction equation in the next iteration step, `jdqr` projects then on the complete search subspace rather than on the current eigenvector approximation.

`Options.LS_Tol` (default [1,0.7,0.7²,...])

`LS_Tol` sets the residual reduction for the linear solver of the correction equation.

If `LS_Tol` is a positive real then the correction equation is solved with a residual reduction of `LS_Tol` in each Jacobi-Davidson step.

If `LS_Tol` is a row $[a_1, a_2, \dots, a_p]$ of positive reals then the correction equation at the i th Jacobi-Davidson step is solved with a residual reduction of a_i provided that $i \leq p$. If $i = p + j$, then the residual reduction at the i th Jacobi-Davidson step is $a_p * b^j$ where $b = a_p / a_{p-1}$.

i is reset to 0 when a Schur vector is detected.

The required residual reduction is not obtained if the maximum number `LS_MaxIt` of iteration steps of the linear solver is reached before.

The default for `LS_Tol` is [1,0.7] if a preconditioner is used (then $a_i = 0.7^{i-1}$). In the other case, the default is [0.7,0.49] (then $a_i = 0.7^i$).

Accuracy

`Options.Tol` (default 1.0e-8)

`jdqr` accepts an approximate Schur vector `q` with associated eigenvalue `lambda` if the 2-norm of the residual `r` is less than `Options.Tol`. Then we have that

$$\text{norm}(A*Q-Q*S, 'fro') < \text{Tol}$$

How accurate the approximations of the eigenvectors and the invariant subspaces are, depends on the conditioning of these quantities.

The residual that `jdqr` computes is given by $\mathbf{r} = (\mathbf{I} - \mathbf{Q} * \mathbf{Q}') * (\mathbf{A} * \mathbf{q} - \text{theta} * \mathbf{B} * \mathbf{q})$.