

# NAG Toolbox for MATLAB

## e04dg

### 1 Purpose

e04dg minimizes an unconstrained nonlinear function of several variables using a pre-conditioned, limited memory quasi-Newton conjugate gradient method. First derivatives (or an ‘acceptable’ finite difference approximation to them) are required. It is intended for use on large scale problems.

### 2 Syntax

```
[iter, objf, objgrd, x, user, lwsav, iwsav, rwsav, ifail] =  
e04dg(objfun, x, lwsav, iwsav, rwsav, 'n', n, 'user', user)
```

Before calling e04dg, or the option setting function e04dk, function e04wb **must** be called.

### 3 Description

e04dg is designed to solve unconstrained minimization problems of the form

$$\underset{x \in R^n}{\text{minimize}} F(x) \quad \text{subject to} \quad -\infty \leq x \leq \infty,$$

where  $x$  is an  $n$  element vector.

You must supply an initial estimate of the solution.

For maximum reliability, it is preferable to provide all first partial derivatives. If all of the derivatives cannot be provided, you are recommended to obtain approximate values (using finite differences) by calling e04xa from within user-supplied (sub)program **objfun**.

The method used by e04dg is described in Section 10.

### 4 References

Gill P E and Murray W 1979 Conjugate-gradient methods for large-scale nonlinear optimization *Technical Report SOL 79-15* Department of Operations Research, Stanford University

Gill P E, Murray W and Wright M H 1981 *Practical Optimization* Academic Press

### 5 Parameters

#### 5.1 Compulsory Input Parameters

1: **objfun** – string containing name of m-file

Its specification is:

```
[mode, objf, objgrd, user] = objfun(mode, n, x, nstate, user)
```

#### Input Parameters

1: **mode** – int32 scalar

Indicates which values must be assigned during each call of **objfun**. Only the following values need be assigned:

**mode** = 0

**objf**.

**mode** = 2

**objf** and **objgrd**.

May be set to a negative value if you wish to terminate the solution to the current problem, and in this case e04dg will terminate with **ifail** set to **mode**.

2: **n** – **int32 scalar**

$n$ , the number of variables.

3: **x(n)** – **double array**

$x$ , the vector of variables at which the objective function and its gradient are to be evaluated.

4: **nstate** – **int32 scalar**

Will be 1 on the first call of **objfun** by e04dg, and 0 for all subsequent calls. Thus, you may wish to test, **nstate** within **objfun** in order to perform certain calculations once only. For example, you may read data or initialize global variables when **nstate** = 1.

5: **user** – Any MATLAB object

**objfun** is called from e04dg with **user** as supplied to e04dg

#### Output Parameters

1: **mode** – **int32 scalar**

Indicates which values must be assigned during each call of **objfun**. Only the following values need be assigned:

**mode** = 0

**objf**.

**mode** = 2

**objf** and **objgrd**.

May be set to a negative value if you wish to terminate the solution to the current problem, and in this case e04dg will terminate with **ifail** set to **mode**.

2: **objf** – **double scalar**

The value of the objective function at  $x$ .

3: **objgrd(n)** – **double array**

If **mode** = 2, **objgrd**( $i$ ) must contain the value of  $\frac{\partial F}{\partial x_i}$  evaluated at  $x$ , for  $i = 1, 2, \dots, n$ .

4: **user** – Any MATLAB object

**objfun** is called from e04dg with **user** as supplied to e04dg

**Note:** **objfun** should be tested separately before being used in conjunction with e04dg. See also the description of the optional parameter **Verify**.

2: **x(n)** – **double array**

An initial estimate of the solution.

- 3: **lwsav(120)** – logical array
- 4: **iwsav(610)** – int32 array
- 5: **rwsav(475)** – double array

The arrays **lwsav**, **iwsav** and **rwsav** **must not** be altered between calls to any of the functions e04wb, e04dg or e04dk.

## 5.2 Optional Input Parameters

- 1: **n** – int32 scalar

$n$ , the number of variables.

*Constraint:*  $n > 0$ .

- 2: **user** – Any MATLAB object

**user** is not used by e04dg, but is passed to **objfun**. Note that for large objects it may be more efficient to use a global variable which is accessible from the m-files than to use **user**.

## 5.3 Input Parameters Omitted from the MATLAB Interface

iwork, work

## 5.4 Output Parameters

- 1: **iter** – int32 scalar

The total number of iterations performed.

- 2: **objf** – double scalar

The value of the objective function at the final iterate.

- 3: **objgrd(n)** – double array

The gradient of the objective function at the final iterate (or its finite difference approximation).

- 4: **x(n)** – double array

The final estimate of the solution.

- 5: **user** – Any MATLAB object

**user** is not used by e04dg, but is passed to **objfun**. Note that for large objects it may be more efficient to use a global variable which is accessible from the m-files than to use **user**.

- 6: **lwsav(120)** – logical array

- 7: **iwsav(610)** – int32 array

- 8: **rwsav(475)** – double array

The arrays **lwsav**, **iwsav** and **rwsav** **must not** be altered between calls to any of the functions e04wb, e04dg or e04dk.

- 9: **ifail** – int32 scalar

0 unless the function detects an error (see Section 6).

## 6 Error Indicators and Warnings

**Note:** e04dg may return useful information for one or more of the following detected errors or warnings.

**ifail** < 0

A negative value of **ifail** indicates an exit from e04dg because you set **mode** < 0 in user-supplied (sub)program **objfun**. The value of **ifail** will be the same as your setting of **mode**.

**ifail** = 1

Not used by this function.

**ifail** = 2

Not used by this function.

**ifail** = 3

The limiting number of iterations (as determined by the optional parameter **Iteration Limit** (default value =  $\max(50, 5n)$ ) has been reached.

If the algorithm appears to be making satisfactory progress, then optional parameter **Iteration Limit** may be too small. If so, increase its value and rerun e04dg. If the algorithm seems to be making little or no progress, then you should check for incorrect gradients as described under **ifail** = 7.

**ifail** = 4

The computed upper bound on the step length taken during the linesearch was too small. A rerun with an increased value of the optional parameter **Maximum Step Length** ( $\rho$  say) may be successful unless  $\rho \geq 10^{20}$  (the default value), in which case the current point cannot be improved upon.

**ifail** = 5

Not used by this function.

**ifail** = 6

The conditions for an acceptable solution (see parameter **ifail** in Section 5) have not all been met, but a lower point could not be found.

If user-supplied (sub)program **objfun** computes the objective function and its gradient correctly, then this may occur because an overly stringent accuracy has been requested, i.e., the value of the optional parameter **Optimality Tolerance** (default value =  $\epsilon^{0.8}$ ) is too small or if  $\alpha_k \simeq 0$ . In this case you should apply the three tests described under **ifail** = 0 to determine whether or not the final solution is acceptable. For a discussion of attainable accuracy see Gill *et al.* 1981.

If many iterations have occurred in which essentially no progress has been made or e04dg has failed to move from the initial point, user-supplied (sub)program **objfun** may be incorrect. You should refer to the comments below under **ifail** = 7 and check the gradients using the optional parameter **Verify** (default value = 0). Unfortunately, there may be small errors in the objective gradients that cannot be detected by the verification process. Finite-difference approximations to first derivatives are catastrophically affected by even small inaccuracies.

**ifail** = 7

The user-supplied derivatives of the objective function appear to be incorrect.

Large errors were found in the derivatives of the objective function. This value of **ifail** will occur if the verification process indicated that at least one gradient element had no correct figures. You should refer to the printed output to determine which elements are suspected to be in error.

As a first step, you should check that the code for the objective values is correct – for example, by computing the function at a point where the correct value is known. However, care should be taken that the chosen point fully tests the evaluation of the function. It is remarkable how often the values  $x = 0$  or  $x = 1$  are used to test function evaluation procedures, and how often the special properties of these numbers make the test meaningless.

Special care should be used in this test if computation of the objective function involves subsidiary data communicated in global storage. Although the first evaluation of the function may be correct, subsequent calculations may be in error because some of the subsidiary data has accidentally been overwritten.

Errors in programming the function may be quite subtle in that the function value is ‘almost’ correct. For example, the function may not be accurate to full precision because of the inaccurate calculation of a subsidiary quantity, or the limited accuracy of data upon which the function depends. A common error on machines where numerical calculations are usually performed in double precision is to include even one single precision constant in the calculation of the function; since some compilers do not convert such constants to double precision, half the correct figures may be lost by such a seemingly trivial error.

**ifail** = 8

The gradient  $\left(g = \frac{\partial F}{\partial x}\right)$  at the starting point  $x_0$  is ‘too small’. More precisely, the value of  $g(x_0)^T g(x_0)$  is less than  $\epsilon_r |1 + F(x_0)|$ , where  $\epsilon_r$  is the value of the optional parameter **Function Precision** (default value =  $\epsilon^{0.9}$ ).

The problem should be rerun from a different starting point.

**ifail** = 9

An input parameter is invalid.

## 7 Accuracy

On successful exit (**ifail** = 0) the accuracy of the solution will be as defined by the optional parameter **Optimality Tolerance** (default value =  $\epsilon^{0.8}$ ).

## 8 Further Comments

To evaluate an ‘acceptable’ set of finite difference intervals using e04xa requires 2 function evaluations per variable for a well-scaled problem and up to 6 function evaluations per variable for a badly scaled problem.

### 8.1 Description of Printed Output

This section describes the intermediate printout and final printout produced by e04dg. You can control the level of printed output (see the description of the optional parameter **Print Level**). Note that the intermediate printout and final printout are produced only if **Print Level**  $\geq 10$  (the default for e04dg, by default no output is produced by e04dg).

The following line of summary output (< 80 characters) is produced at every iteration. In all cases, the values of the quantities are those in effect *on completion* of the given iteration.

<b>Itn</b>	is the iteration count.
<b>Step</b>	is the step $\alpha_k$ taken along the computed search direction. On reasonably well-behaved problems, the unit step (i.e., $\alpha_k = 1$ ) will be taken as the solution is approached.
<b>Nfun</b>	is the cumulated number of evaluations of the objective function needed for the linesearch. Evaluations needed for the verification of the gradients by finite differences are not included. Nfun is printed as a guide to the amount of work required for the linesearch. e04dg will perform at most 11 function evaluations per iteration.
<b>Objective</b>	is the value of the objective function at $x_k$ .
<b>Norm G</b>	is the Euclidean norm of the gradient of the objective function at $x_k$ .

Norm  $X$  is the Euclidean norm of  $x_k$ .

Norm  $(X(k-1)-X(k))$  is the Euclidean norm of  $x_{k-1} - x_k$ .

The following describes the printout for each variable.

Variable gives the name (Varbl) and index  $j$ , for  $j = 1, 2, \dots, n$  of the variable.

Value is the value of the variable at the final iteration.

Gradient Value is the value of the gradient of the objective function with respect to the  $j$ th variable at the final iteration.

Numerical values are output with a fixed number of digits; they are not guaranteed to be accurate to this precision.

## 9 Example

```
e04dg_objfun.m
```

```
function [mode, objf, objgrd, user] = e04dg_objfun(mode, n, x, nstate,
user)

    expx1 = exp(x(1));
    objf = expx1*(4.0*x(1)^2+2.0*x(2)^2+4.0*x(1)*x(2)+2.0*x(2)+1.0);
    if (mode == 2)
        objgrd(1) = 4.0*expx1*(2.0*x(1)+x(2)) + objf;
        objgrd(2) = 2.0*expx1*(2.0*x(2)+2.0*x(1)+1.0);
    else
        objgrd = zeros(2,1);
    end
```

```
x = [-1;
      1];
[cwsav,lwsav,iwsav,rwsav,ifail] = e04wb('e04dg');
[iter, objf, objgrd, xOut, user, lwsavOut, iwsavOut, rwsavOut, ifail] =
...
    e04dg('e04dg_objfun', x, lwsav, iwsav, rwsav)
```

```
iter =
      10
objf =
  5.3083e-14
objgrd =
  1.0e-06 *
    0.9125
    0.8316
xOut =
    0.5000
   -1.0000
user =
    0
lwsavOut =
    array elided
iwsavOut =
    array elided
rwsavOut =
    array elided
ifail =
    0
```

**Note:** the remainder of this document is intended for more advanced users. Section 10 contains a detailed description of the algorithm which may be needed in order to understand Section 11. Section 11 describes the optional parameters which may be set by calls to e04dk.

## 10 Algorithmic Details

This section contains a description of the method used by e04dg.

e04dg uses a pre-conditioned conjugate gradient method and is based upon algorithm PLMA as described in Section 4.8.3 of Gill and Murray 1979 and Gill *et al.* 1981.

The algorithm proceeds as follows:

Let  $x_0$  be a given starting point and let  $k$  denote the current iteration, starting with  $k = 0$ . The iteration requires  $g_k$ , the gradient vector evaluated at  $x_k$ , the  $k$ th estimate of the minimum. At each iteration a vector  $p_k$  (known as the direction of search) is computed and the new estimate  $x_{k+1}$  is given by  $x_k + \alpha_k p_k$  where  $\alpha_k$  (the step length) minimizes the function  $F(x_k + \alpha_k p_k)$  with respect to the scalar  $\alpha_k$ . A choice of initial step  $\alpha_0$  is taken as

$$\alpha_0 = \min\{1, 2 \times |F_k - F_{\text{est}}| / g_k^T g_k\}$$

where  $F_{\text{est}}$  is a user-supplied estimate of the function value at the solution. If  $F_{\text{est}}$  is not specified, the software always chooses the unit step length for  $\alpha_0$ . Subsequent step length estimates are computed using cubic interpolation with safeguards.

A quasi-Newton method can be used to compute the search direction  $p_k$  by updating the inverse of the approximate Hessian ( $H_k$ ) and computing

$$p_{k+1} = -H_{k+1} g_{k+1}. \quad (1)$$

The updating formula for the approximate inverse is given by

$$H_{k+1} = H_k - \frac{1}{y_k^T s_k} (H_k y_k s_k^T + s_k y_k^T H_k) + \frac{1}{y_k^T s_k} \left( 1 + \frac{y_k^T H_k y_k}{y_k^T s_k} \right) s_k s_k^T, \quad (2)$$

where  $y_k = g_{k+1} - g_k$  and  $s_k = x_{k+1} - x_k = \alpha_k p_k$ .

The method used to obtain the search direction is based upon computing  $p_{k+1}$  as  $-H_{k+1} g_{k+1}$  where  $H_{k+1}$  is a matrix obtained by updating the identity matrix with a limited number of quasi-Newton corrections. The storage of an  $n$  by  $n$  matrix is avoided by storing only the vectors that define the rank two corrections – hence the term ‘limited-memory’ quasi-Newton method. The precise method depends upon the number of updating vectors stored. For example, the direction obtained with the ‘one-step’ limited memory update is given by (1) using (2) with  $H_k$  equal to the identity matrix, viz.

$$p_{k+1} = -g_{k+1} + \frac{1}{y_k^T s_k} (s_k^T g_{k+1} y_k + y_k^T g_{k+1} s_k) - \frac{s_k^T g_{k+1}}{y_k^T s_k} \left( 1 + \frac{y_k^T y_k}{y_k^T s_k} \right) s_k.$$

Using a limited-memory quasi-Newton formula, such as the one above, guarantees  $p_{k+1}$  to be a descent direction if all the inner products  $y_k^T s_k$  are positive for all vectors  $y_k$  and  $s_k$  used in the updating formula.

## 11 Optional Parameters

Several optional parameters in e04dg define choices in the problem specification or the algorithm logic. In order to reduce the number of formal parameters of e04dg these optional parameters have associated *default values* that are appropriate for most problems. Therefore, you need only specify those optional parameters whose values are to be different from their default values.

The remainder of this section can be skipped if you wish to use the default values for *all* optional parameters. A complete list of optional parameters and their default values is given in Section 11.1.

Optional parameters may be specified by calling e04dk before a call to e04dg.

e04dk can be called to supply options directly, one call being necessary for each optional parameter. For example,

```
[lwsav, iwsav, rwsav, inform] = e04dk('Print Level = 1', lwsav, iwsav, rwsav);
```

e04dk should be consulted for a full description of this method of supplying optional parameters.

All optional parameters not specified by you are set to their default values. Optional parameters specified by you are unaltered by e04dg (unless they define invalid values) and so remain in effect for subsequent calls unless altered by you.

## 11.1 Optional Parameter Checklist and Default Values

The following list gives the valid options. For each option, we give the keyword, any essential optional qualifiers and the default value. A definition for each option can be found in Section 11.2. The minimum abbreviation of each keyword is underlined. If no characters of an optional qualifier are underlined, the qualifier may be omitted. The letter *a* denotes a phrase (character string) that qualifies an option. The letters *i* and *r* denote integer and double values required with certain options. The number  $\epsilon$  is a generic notation for *machine precision* (see x02aj), and  $\epsilon_r$  denotes the relative precision of the objective function.

Optional Parameters	Default Values
<b><u>Defaults</u></b>	
<b><u>Estimated Optimal Function Value</u></b>	
<b><u>Function Precision</u></b>	Default = $\epsilon^{0.9}$
<b><u>Iteration Limit</u></b>	Default = $\max(50, 5n)$
<b><u>Iters</u></b>	See <b><u>Iteration Limit</u></b>
<b><u>Itns</u></b>	See <b><u>Iteration Limit</u></b>
<b><u>Linesearch Tolerance</u></b>	Default = 0.9
<b><u>List</u></b>	Default for e04dg = <b>List</b>
<b><u>Maximum Step Length</u></b>	Default = $10^{20}$
<b><u>Nolist</u></b>	Default for e04dg = <b>Nolist</b> . See <b><u>List</u></b> .
<b><u>Optimality Tolerance</u></b>	Default = $\epsilon_R^{0.8}$
<b><u>Print Level</u></b>	Default for e04dg = 10 Default for e04dg = 0
<b><u>Start Objective Check at Variable</u></b>	Default = 1
<b><u>Stop Objective Check at Variable</u></b>	Default = <i>n</i> . See <b><u>Start Objective Check at Variable</u></b> .
<b><u>Verify</u></b>	See <b><u>Verify Level</u></b>
<b><u>Verify Gradients</u></b>	See <b><u>Verify Level</u></b>
<b><u>Verify Level</u></b>	Default = 0
<b><u>Verify Objective Gradients</u></b>	See <b><u>Verify Level</u></b>

## 11.2 Description of the Optional Parameters

### Defaults

This special keyword may be used to reset all optional parameters to their default values.

### Estimated Optimal Function Value *r*

This value of *r* specifies the user-supplied guess of the optimum objective function value  $F_{\text{est}}$ . This value is used to calculate an initial step length  $\alpha_0$  (see Section 10). If the value of *r* is not specified (the default), then this has the effect of setting  $\alpha_0$  to unity. It should be noted that for badly scaled functions a unit step along the steepest descent direction will often compute the objective function at very large values of *x*.

### Function Precision *r* Default = $\epsilon^{0.9}$

The parameter defines  $\epsilon_r$ , which is intended to be a measure of the accuracy with which the problem function  $F(x)$  can be computed. If  $r < \epsilon$  or  $r \geq 1$ , the default value is used.



The value of  $\epsilon_r$  should reflect the relative precision of  $1 + |F(x)|$ ; i.e.,  $\epsilon_r$  acts as a relative precision when  $|F|$  is large, and as an absolute precision when  $|F|$  is small. For example, if  $F(x)$  is typically of order 1000 and the first six significant digits are known to be correct, an appropriate value for  $\epsilon_r$  would be  $10^{-6}$ . In contrast, if  $F(x)$  is typically of order  $10^{-4}$  and the first six significant digits are known to be correct, an appropriate value for  $\epsilon_r$  would be  $10^{-10}$ . The choice of  $\epsilon_r$  can be quite complicated for badly scaled problems; see Chapter 8 of Gill *et al.* 1981 for a discussion of scaling techniques. The default value is appropriate for most simple functions that are computed with full accuracy. However when the accuracy of the computed function values is known to be significantly worse than full precision, the value of  $\epsilon_r$  should be large enough so that no attempt will be made to distinguish between function values that differ by less than the error inherent in the calculation.

<b><u>Iteration Limit</u></b>	$i$	Default = max(50, 5n)
<b><u>Iters</u></b>		
<b><u>Itns</u></b>		

The value of  $i$  specifies the maximum number of iterations allowed before termination. If  $i < 0$ , the default value is used.

Problems whose Hessian matrices at the solution contain sets of clustered eigenvalues are likely to be minimized in significantly fewer than  $n$  iterations. Problems without this property may require anything between  $n$  and  $5n$  iterations, with approximately  $2n$  iterations being a common figure for moderately difficult problems.

<b><u>Linesearch Tolerance</u></b>	$r$	Default = 0.9
------------------------------------	-----	---------------

The value  $r$  controls the accuracy with which the step  $\alpha$  taken during each iteration approximates a minimum of the function along the search direction (the smaller the value of  $r$ , the more accurate the linesearch). The default value  $r = 0.9$  requests an inaccurate search, and is appropriate for most problems. A more accurate search may be appropriate when it is desirable to reduce the number of iterations – for example, if the objective function is cheap to evaluate. If  $r < 0$  or  $r \geq 1$ , the default value is used.

<b><u>List</u></b>	Default for e04dg = <b>List</b>
<b><u>Nolist</u></b>	Default for e04dg = <b>Nolist</b>

Normally each optional parameter specification is printed as it is supplied. Optional parameter **Nolist** may be used to suppress the printing and optional parameter **List** may be used to restore printing.

<b><u>Maximum Step Length</u></b>	$r$	Default = $10^{20}$
-----------------------------------	-----	---------------------

If  $r > 0$ , the maximum allowable step length for the linesearch is taken as  $\min(1 \times 10^{20}, \frac{r}{\|p_k\|})$ . If  $r \leq 0$ , the default value is used.

<b><u>Optimality Tolerance</u></b>	$r$	Default = $\epsilon_R^{0.8}$
------------------------------------	-----	------------------------------

The parameter  $r$  specifies the accuracy to which you wish the final iterate to approximate a solution of the problem. Broadly speaking,  $r$  indicates the number of correct figures desired in the objective function at the solution. For example, if  $r$  is  $10^{-6}$  and termination occurs with **ifail** = 0 (see Section 5), then the final point satisfies the termination criteria, where  $\tau_F$  represents **Optimality Tolerance**. If  $r < \epsilon_r$  or  $r \geq 1$ , the default value is used.

<b><u>Print Level</u></b>	$i$	Default for e04dg = 10 Default for e04dg = 0
---------------------------	-----	---

The value  $i$  controls the amount of printout produced by e04dg, as indicated below. A detailed description of the printout is given in Section 8.1 (summary output at each iteration and the final solution).

$i$	<b>Output</b>
0	No output.

- 1 The final solution only.
- 5 One line of summary output ( < 80 characters; see Section 8.1) for each iteration (no printout of the final solution).
- 10 The final solution and one line of summary output for each iteration.

<b>Start Objective Check at Variable</b>	$i_1$	Default = 1
<b>Stop Objective Check at Variable</b>	$i_2$	Default = $n$

These keywords take effect only if **Verify Level** > 0. They may be used to control the verification of gradient elements computed by user-supplied (sub)program **objfun**. For example, if the first 30 elements of the objective gradient appeared to be correct in an earlier run, so that only element 31 remains questionable, it is reasonable to specify **Start Objective Check at Variable** = 31. If the first 30 variables appear linearly in the objective, so that the corresponding gradient elements are constant, the above choice would also be appropriate.

If  $i_1 \leq 0$  or  $i_1 > \max(1, \min(n, i_2))$ , the default value is used. If  $i_2 \leq 0$  or  $i_2 > n$ , the default value is used.

<b>Verify Level</b>	$i$	Default = 0
---------------------	-----	-------------

**Verify**  
**Verify Gradients**  
**Verify Objective Gradients**

These keywords refer to finite-difference checks on the gradient elements computed by user-supplied (sub)program **objfun**. Gradients are verified at the user-supplied initial estimate of the solution. The possible choices for  $i$  are as follows:

$i$	Meaning
-1	No checks are performed.
0	Only a 'cheap' test will be performed, requiring one call to user-supplied (sub)program <b>objfun</b> .
1	In addition to the 'cheap' test, individual gradient elements will also be checked using a reliable (but more expensive) test.

For example, the objective gradient will be verified if **Verify**, **Verify** = Yes, **Verify Gradients**, **Verify Objective Gradients** or **Verify Level** = 1 is specified.