

NAG Toolbox

nag_glopt_bnd_pso (e05sa)

1 Purpose

nag_glopt_bnd_pso (e05sa) is designed to search for the global minimum or maximum of an arbitrary function, using Particle Swarm Optimization (PSO). Derivatives are not required, although these may be used by an accompanying local minimization function if desired. nag_glopt_bnd_pso (e05sa) is essentially identical to nag_glopt_nlp_pso (e05sb), but with a simpler interface and with various optional parameters removed; otherwise most arguments are identical. In particular, nag_glopt_bnd_pso (e05sa) does not handle general constraints.

2 Syntax

```
[xb, fb, iopts, opts, user, itt, inform, ifail] = nag_glopt_bnd_pso(bl, bu,
objfun, monmod, iopts, opts, 'ndim', ndim, 'npar', npar, 'user', user)
```

```
[xb, fb, iopts, opts, user, itt, inform, ifail] = e05sa(bl, bu, objfun, monmod,
iopts, opts, 'ndim', ndim, 'npar', npar, 'user', user)
```

Before calling nag_glopt_bnd_pso (e05sa), nag_glopt_optset (e05zk) **must** be called with **optstr** set to ‘’. Optional parameters may also be specified by calling nag_glopt_optset (e05zk) before the call to nag_glopt_bnd_pso (e05sa).

Note: the interface to this routine has changed since earlier releases of the toolbox:

At Mark 24: **npar** was made optional.

3 Description

nag_glopt_bnd_pso (e05sa) uses a stochastic method based on Particle Swarm Optimization (PSO) to search for the global optimum of a nonlinear function F , subject to a set of bound constraints on the variables. In the PSO algorithm (see Section 11), a set of particles is generated in the search space, and advances each iteration to (hopefully) better positions using a heuristic velocity based upon *inertia*, *cognitive memory* and *global memory*. The inertia is provided by a decreasingly weighted contribution from a particles current velocity, the cognitive memory refers to the best candidate found by an individual particle and the global memory refers to the best candidate found by all the particles. This allows for a global search of the domain in question.

Further, this may be coupled with a selection of local minimization functions, which may be called during the iterations of the heuristic algorithm, the *interior* phase, to hasten the discovery of locally optimal points, and after the heuristic phase has completed to attempt to refine the final solution, the *exterior* phase. Different options may be set for the local optimizer in each phase.

Without loss of generality, the problem is assumed to be stated in the following form:

$$\underset{\mathbf{x} \in R^{ndim}}{\text{minimize}} F(\mathbf{x}) \quad \text{subject to} \quad \boldsymbol{\ell} \leq \mathbf{x} \leq \mathbf{u},$$

where the objective $F(\mathbf{x})$ is a scalar function, \mathbf{x} is a vector in R^{ndim} and the vectors $\boldsymbol{\ell} \leq \mathbf{u}$ are lower and upper bounds respectively for the $ndim$ variables. The objective function may be nonlinear. Continuity of F is not essential. For functions which are smooth and primarily unimodal, faster solutions will almost certainly be achieved by using Chapter E04 functions directly.

For functions which are smooth and multi-modal, gradient dependent local minimization functions may be coupled with nag_glopt_bnd_pso (e05sa).

For multi-modal functions for which derivatives cannot be provided, particularly functions with a significant level of noise in their evaluation, `nag_glopt_bnd_pso` (e05sa) should be used either alone, or coupled with `nag_opt_uncon_simplex` (e04cb).

The *ndim* lower and upper box bounds on the variable **x** are included to initialize the particle swarm into a finite hypervolume, although their subsequent influence on the algorithm is user determinable (see the option **Boundary** in Section 12). It is strongly recommended that sensible bounds are provided for all variables.

`nag_glopt_bnd_pso` (e05sa) may also be used to maximize the objective function (see the option **Optimize**).

Due to the nature of global optimization, unless a predefined target is provided, there is no definitive way of knowing when to end a computation. As such several stopping heuristics have been implemented into the algorithm. If any of these is achieved, `nag_glopt_bnd_pso` (e05sa) will exit with **ifail** = 1, and the parameter **inform** will indicate which criteria was reached. See **inform** for more information.

In addition, you may provide your own stopping criteria through **monmod** and **objfun**.

`nag_glopt_nlp_pso` (e05sb) provides a comprehensive interface, allowing for the inclusion of general nonlinear constraints.

4 References

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

Kennedy J and Eberhart R C (1995) Particle Swarm Optimization *Proceedings of the 1995 IEEE International Conference on Neural Networks 1942–1948*

Koh B, George A D, Haftka R T and Fregly B J (2006) Parallel Asynchronous Particle Swarm Optimization *International Journal for Numerical Methods in Engineering* **67(4)** 578–595

Vaz A I and Vicente L N (2007) A Particle Swarm Pattern Search Method for Bound Constrained Global Optimization *Journal of Global Optimization* **39(2)** 197–219 Kluwer Academic Publishers

5 Parameters

Note: for descriptions of the symbolic variables, see Section 11.

5.1 Compulsory Input Parameters

- 1: **bl(ndim)** – REAL (KIND=nag_wp) array
- 2: **bu(ndim)** – REAL (KIND=nag_wp) array

bl is **ℓ**, the array of lower bounds, **bu** is **u**, the array of upper bounds. The **ndim** entries in **bl** and **bu** must contain the lower and upper simple (box) bounds of the variables respectively. These must be provided to initialize the sample population into a finite hypervolume, although their subsequent influence on the algorithm is user determinable (see the option **Boundary** in Section 12).

If **bl**(*i*) = **bu**(*i*) for any $i \in \{1, \dots, \mathbf{ndim}\}$, variable *i* will remain locked to **bl**(*i*) regardless of the **Boundary** option selected.

It is strongly advised that you place sensible lower and upper bounds on all variables, even if your model allows for variables to be unbounded (using the option **Boundary** = ignore) since these define the initial search space.

Constraints:

$$\begin{aligned} \mathbf{bl}(i) &\leq \mathbf{bu}(i), \text{ for } i = 1, 2, \dots, \mathbf{ndim}; \\ \mathbf{bl}(i) &\neq \mathbf{bu}(i) \text{ for at least one } i \in \{1, \dots, \mathbf{ndim}\}. \end{aligned}$$

3: **objfun** – SUBROUTINE, supplied by the user.

objfun must, depending on the value of **mode**, calculate the objective function *and/or* calculate the gradient of the objective function for a *ndim*-variable vector **x**. Gradients are only required if a local minimizer has been chosen which requires gradients. See the option **Local Minimizer** for more information.

```
[mode, objf, vecout, user] = objfun(mode, ndim, x, objf, vecout, nstate, user)
```

Input Parameters

1: **mode** – INTEGER

Indicates which functionality is required.

mode = 0

$F(\mathbf{x})$ should be returned in **objf**. The value of **objf** on entry may be used as an upper bound for the calculation. Any expected value of $F(\mathbf{x})$ that is greater than **objf** may be approximated by this upper bound; that is **objf** can remain unaltered.

mode = 1

Local Minimizer = e04uc only

First derivatives can be evaluated and returned in **vecout**. Any unaltered elements of **vecout** will be approximated using finite differences.

mode = 2

Local Minimizer = e04uc only

$F(\mathbf{x})$ *must* be calculated and returned in **objf**, and available first derivatives can be evaluated and returned in **vecout**. Any unaltered elements of **vecout** will be approximated using finite differences.

mode = 5

$F(\mathbf{x})$ *must* be calculated and returned in **objf**. The value of **objf** on entry may not be used as an upper bound.

mode = 6

Local Minimizer = e04dg or e04kz only

All first derivatives *must* be evaluated and returned in **vecout**.

mode = 7

Local Minimizer = e04dg or e04kz only

$F(\mathbf{x})$ *must* be calculated and returned in **objf**, and *all* first derivatives *must* be evaluated and returned in **vecout**.

2: **ndim** – INTEGER

The number of dimensions.

3: **x(ndim)** – REAL (KIND=nag_wp) array

x, the point at which the objective function and/or its gradient are to be evaluated.

4: **objf** – REAL (KIND=nag_wp)

The value of **objf** passed to **objfun** varies with the argument **mode**.

mode = 0

objf is an upper bound for the value of $F(\mathbf{x})$, often equal to the best value of $F(\mathbf{x})$ found so far by a given particle. Only objective function values less than the value of **objf** on entry will be used further. As such this upper bound may be used to stop further evaluation when this will only increase the objective function value above the upper bound.

mode = 1, 2, 5, 6 or 7
objf is meaningless on entry.

5: **vecout**(**ndim**) – REAL (KIND=nag_wp) array

If **Local Minimizer** = E04UCF or E04UCA, the values of **vecout** are used internally to indicate whether a finite difference approximation is required. See nag_opt_nlp1_solve (e04uc).

6: **nstate** – INTEGER

nstate indicates various stages of initialization throughout the function. This allows for permanent global arguments to be initialized the least number of times. For example, you may initialize a random number generator seed.

nstate = 3

SMP users only. **objfun** is called for the first time in a parallel region on a new thread other than the master thread. You may use this opportunity to set up any thread-dependent information in **user** and **user**.

nstate = 2

objfun is called for the very first time. You may save computational time if certain data must be read or calculated only once.

nstate = 1

objfun is called for the first time by a NAG local minimization function. You may save computational time if certain data required for the local minimizer need only be calculated at the initial point of the local minimization.

nstate = 0

Used in all other cases.

7: **user** – INTEGER array

objfun is called from nag_glopt_bnd_pso (e05sa) with the object supplied to nag_glopt_bnd_pso (e05sa).

Output Parameters

1: **mode** – INTEGER

If the value of **mode** is set to be negative, then nag_glopt_bnd_pso (e05sa) will exit as soon as possible with **ifail** = 3 and **inform** = **mode**.

2: **objf** – REAL (KIND=nag_wp)

The value of **objf** returned varies with the argument **mode**.

mode = 0

objf must be the value of $F(\mathbf{x})$. Only values of $F(\mathbf{x})$ strictly less than **objf** on entry need be accurate.

mode = 1 or 6

Need not be set.

mode = 2, 5 or 7

$F(\mathbf{x})$ must be calculated and returned in **objf**. The entry value of **objf** may not be used as an upper bound.

3: **vecout**(**ndim**) – REAL (KIND=nag_wp) array

The required values of **vecout** returned to the calling function depend on the value of **mode**.

mode = 0 or 5
The value of **vecout** need not be set.

mode = 1 or 2
vecout can contain components of the gradient of the objective function $\frac{\partial F}{\partial x_i}$ for some $i = 1, 2, \dots, \mathbf{ndim}$, or acceptable approximations. Any unaltered elements of **vecout** will be approximated using finite differences.

mode = 6 or 7
vecout must contain the gradient of the objective function $\frac{\partial F}{\partial x_i}$ for all $i = 1, 2, \dots, \mathbf{ndim}$. Approximation of the gradient is strongly discouraged, and no finite difference approximations will be performed internally (see nag_opt_uncon_conjgrd_comp (e04dg) and nag_opt_bounds_mod_deriv_easy (e04kz)).

4: **user** – INTEGER array

4: **monmod** – SUBROUTINE, supplied by the NAG Library or the user.

A user-specified monitoring and modification function. **monmod** is called once every complete iteration after a finalization check. It may be used to modify the particle locations that will be evaluated at the next iteration. This permits the incorporation of algorithmic modifications such as including additional advection heuristics and genetic mutations. **monmod** is only called during the main loop of the algorithm, and as such will be unaware of any further improvement from the final local minimization. If no monitoring and/or modification is required, **monmod** may be string nag_glopt_bnd_pso_dummy_monmod (e05sxm) nag_glopt_bnd_pso_dummy_monmod (e05sxm) .

```
[x, user, inform] = monmod(ndim, npar, x, xb, fb, xbest, fbest, itt, user,
inform)
```

Input Parameters

- 1: **ndim** – INTEGER
The number of dimensions.
- 2: **npar** – INTEGER
The number of particles.
- 3: **x**(**ndim**, **npar**) – REAL (KIND=nag_wp) array
Note: the i th component of the j th particle, $x_j(i)$, is stored in $\mathbf{x}(i, j)$.
The **npar** particle locations, \mathbf{x}_j , which will currently be used during the next iteration unless altered in **monmod**.
- 4: **xb**(**ndim**) – REAL (KIND=nag_wp) array
The location, $\tilde{\mathbf{x}}$, of the best solution yet found.
- 5: **fb** – REAL (KIND=nag_wp)
The objective value, $\tilde{f} = F(\tilde{\mathbf{x}})$, of the best solution yet found.

- 6: **xbest**(**ndim**, **npar**) – REAL (KIND=nag_wp) array
Note: the i th component of the position of the j th particle's cognitive memory, $\hat{x}_j(i)$, is stored in **xbest**(i, j).
The locations currently in the cognitive memory, $\hat{\mathbf{x}}_j$, for $j = 1, 2, \dots, \mathbf{npar}$ (see Section 11).
- 7: **fbest**(**npar**) – REAL (KIND=nag_wp) array
The objective values currently in the cognitive memory, $F(\hat{\mathbf{x}}_j)$, for $j = 1, 2, \dots, \mathbf{npar}$.
- 8: **itt**(6) – INTEGER array
Iteration and function evaluation counters (see description of **itt** below).
- 9: **user** – INTEGER array
monmod is called from `nag_glopt_bnd_pso` (e05sa) with the object supplied to `nag_glopt_bnd_pso` (e05sa).
- 10: **inform** – INTEGER
inform = **thread_num**, where **thread_num** is the value returned by a call of the OpenMP function `OMP_GET_THREAD_NUM()`. If running in serial this will always be zero.

Output Parameters

- 1: **x**(**ndim**, **npar**) – REAL (KIND=nag_wp) array
Note: the i th component of the j th particle, $x_j(i)$, is stored in **x**(i, j).
The particle locations to be used during the next iteration.
- 2: **user** – INTEGER array
- 3: **inform** – INTEGER
Setting **inform** < 0 will cause near immediate exit from `nag_glopt_bnd_pso` (e05sa). This value will be returned as **inform** with **ifail** = 3. You need not set **inform** unless you wish to force an exit.
- 5: **iopts**(:) – INTEGER array
Note: the dimension of this array is dictated by the requirements of associated functions that must have been previously called. This array **must** be the same array passed as argument **iopts** in the previous call to `nag_glopt_optset` (e05zk).
Optional parameter array as generated and possibly modified by calls to `nag_glopt_optset` (e05zk). The contents of **iopts** **must not** be modified directly between calls to `nag_glopt_bnd_pso` (e05sa), `nag_glopt_optset` (e05zk) or `nag_glopt_optget` (e05zl).
- 6: **opts**(:) – REAL (KIND=nag_wp) array
Note: the dimension of this array is dictated by the requirements of associated functions that must have been previously called. This array **must** be the same array passed as argument **opts** in the previous call to `nag_glopt_optset` (e05zk).
Optional parameter array as generated and possibly modified by calls to `nag_glopt_optset` (e05zk). The contents of **opts** **must not** be modified directly between calls to `nag_glopt_bnd_pso` (e05sa), `nag_glopt_optset` (e05zk) or `nag_glopt_optget` (e05zl).

5.2 Optional Input Parameters

1: **ndim** – INTEGER

Default: the dimension of the arrays **bl**, **bu**. (An error is raised if these dimensions are not equal.)
ndim, the number of dimensions.

Constraint: **ndim** \geq 1.

2: **npar** – INTEGER

Default: **npar** = $10 \times$ **ndim**

npar, the number of particles to be used in the swarm. Assuming all particles remain within bounds, each complete iteration will perform at least **npar** function evaluations. Otherwise, significantly fewer objective function evaluations may be performed.

Constraint: **npar** \geq $5 \times$ **num_threads**, where **num_threads** is the value returned by the OpenMP environment variable `OMP_NUM_THREADS`, or **num_threads** is 1 for a serial version of this function.

3: **user** – INTEGER array

user is not used by `nag_glopt_bnd_pso` (e05sa), but is passed to **objfun** and **monmod**. 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 Output Parameters

1: **xb(ndim)** – REAL (KIND=nag_wp) array

The location of the best solution found, $\tilde{\mathbf{x}}$, in R^{ndim} .

2: **fb** – REAL (KIND=nag_wp)

The objective value of the best solution, $\tilde{f} = F(\tilde{\mathbf{x}})$.

3: **iopts(:)** – INTEGER array

Note: the dimension of this array is dictated by the requirements of associated functions that must have been previously called. This array **must** be the same array passed as argument **iopts** in the previous call to `nag_glopt_optset` (e05zk).

Communication array, used to store information between calls to `nag_glopt_bnd_pso` (e05sa).

4: **opts(:)** – REAL (KIND=nag_wp) array

Note: the dimension of this array is dictated by the requirements of associated functions that must have been previously called. This array **must** be the same array passed as argument **opts** in the previous call to `nag_glopt_optset` (e05zk).

Communication array, used to store information between calls to `nag_glopt_bnd_pso` (e05sa).

5: **user** – INTEGER array

6: **itt(6)** – INTEGER array

Integer iteration counters for `nag_glopt_bnd_pso` (e05sa).

itt(1)

Number of complete iterations.

itt(2)

Number of complete iterations without improvement to the current optimum.

- itt(3)** Number of particles converged to the current optimum.
- itt(4)** Number of improvements to the optimum.
- itt(5)** Number of function evaluations performed.
- itt(6)** Number of particles reset.

7: **inform** – INTEGER

Indicates which finalization criterion was reached. The possible values of **inform** are:

inform	Meaning
< 0	Exit from a user-supplied subroutine.
0	e05sa has detected an error and terminated.
1	The provided objective target has been achieved. (Target Objective Value).
2	The standard deviation of the location of all the particles is below the set threshold (Swarm Standard Deviation). If the solution returned is not satisfactory, you may try setting a smaller value of Swarm Standard Deviation , or try adjusting the options governing the repulsive phase (Repulsion Initialize , Repulsion Finalize).
3	The total number of particles converged (Maximum Particles Converged) to the current global optimum has reached the set limit. This is the number of particles which have moved to a distance less than Distance Tolerance from the optimum with regard to the L^2 norm. If the solution is not satisfactory, you may consider lowering the Distance Tolerance . However, this may hinder the global search capability of the algorithm.
4	The maximum number of iterations without improvement (Maximum Iterations Static) has been reached, and the required number of particles (Maximum Iterations Static Particles) have converged to the current optimum. Increasing either of these options will allow the algorithm to continue searching for longer. Alternatively if the solution is not satisfactory, re-starting the application several times with Repeatability = OFF may lead to an improved solution.
5	The maximum number of iterations (Maximum Iterations Completed) has been reached. If the number of iterations since improvement is small, then a better solution may be found by increasing this limit, or by using the option Local Minimizer with corresponding exterior options. Otherwise if the solution is not satisfactory, you may try re-running the application several times with Repeatability = OFF and a lower iteration limit, or adjusting the options governing the repulsive phase (Repulsion Initialize , Repulsion Finalize).
6	The maximum allowed number of function evaluations (Maximum Function Evaluations) has been reached. As with inform = 5, increasing this limit if the number of iterations without improvement is small, or decreasing this limit and running the algorithm multiple times with Repeatability = OFF, may provide a superior result.

8: **ifail** – INTEGER

ifail = 0 unless the function detects an error (see Section 5).

For this reason, the value -1 or 1 is recommended. If the output of error messages is undesirable, then the value 1 is recommended; otherwise, the recommended value is -1 . **When the value -1 or 1 is used it is essential to test the value of `ifail` on exit.**

`nag_glopt_bnd_pso` (e05sa) will return `ifail` = 0 if and only if a finalization criterion has been reached which can guarantee success. This may only happen if the option **Target Objective Value** has been set and reached at a point within the search domain. The finalization criterion **Target Objective Value** is not activated using default option settings, and must be explicitly set using `nag_glopt_optset` (e05zk) if required.

`nag_glopt_bnd_pso` (e05sa) will return `ifail` = 1 if no error has been detected, and a finalization criterion has been achieved which cannot guarantee success. This does not indicate that the function has failed, merely that the returned solution cannot be guaranteed to be the true global optimum.

The value of `inform` should be examined to determine which finalization criterion was reached.

Other positive values of `ifail` indicate that either an error or a warning has been triggered. See Section 6, Section 7 and Section 11 for more information.

6 Error Indicators and Warnings

Errors or warnings detected by the function:

`ifail` = 1 (*warning*)

A finalization criterion was reached that cannot guarantee success.

`ifail` = 2 (*warning*)

If the option **Target Warning** has been activated, this indicates that the **Target Objective Value** has been achieved to specified tolerances at a sufficiently constrained point, either during the initialization phase, or during the first two iterations of the algorithm. While this is not necessarily an error, it may occur if:

- (i) The target was achieved at the first point sampled by the function. This will be the mean of the lower and upper bounds.
- (ii) The target may have been achieved at a randomly generated sample point. This will always be a possibility provided that the domain under investigation contains a point with a target objective value.
- (iii) If the **Local Minimizer** has been set, then a sample point may have been inside the basin of attraction of a satisfactory point. If this occurs repeatedly when the function is called, it may imply that the objective is largely unimodal, and that it may be more efficient to use the function selected as the **Local Minimizer** directly.

Assuming that `objfun` is correct, you may wish to set a better **Target Objective Value**, or a stricter **Target Objective Tolerance**.

`ifail` = 3 (*warning*)

User requested exit during call to `monmod`.

User requested exit during call to `objfun`.

`ifail` = 11

Constraint: `ndim` ≥ 1 .

`ifail` = 12

Constraint: `npar` $\geq 5 \times$ `num_threads`, where `num_threads` is the value returned by the OpenMP environment variable `OMP_NUM_THREADS`, or `num_threads` is 1 for a serial version of this function.

ifail = 14

Constraint: $\mathbf{bu}(i) \geq \mathbf{bl}(i)$ for all i .

On entry, $\mathbf{bl}(i) = \mathbf{bu}(i)$ for all i .

Constraint: $\mathbf{bu}(i) > \mathbf{bl}(i)$ for at least one i .

ifail = 19

Error $\langle value \rangle$ occurred whilst adjusting to exterior local minimizer options.

Error $\langle value \rangle$ occurred whilst adjusting to interior local minimizer options.

ifail = 21

Either the option arrays have not been initialized for `nag_glopt_bnd_pso` (e05sa), or they have become corrupted.

ifail = 32 (*warning*)

Derivative checks indicate possible errors in the supplied derivatives. Gradient checks may be disabled by setting **Verify Gradients** = OFF.

ifail = -99

An unexpected error has been triggered by this routine. Please contact NAG.

ifail = -399

Your licence key may have expired or may not have been installed correctly.

ifail = -999

Dynamic memory allocation failed.

7 Accuracy

If **ifail** = 0 (or **ifail** = 2) or **ifail** = 1 on exit, either a **Target Objective Value** or finalization criterion has been reached, depending on user selected options. As with all global optimization software, the solution achieved may not be the true global optimum. Various options allow for either greater search diversity or faster convergence to a (local) optimum (See Section 11 and Section 12).

Provided the objective function and constraints are sufficiently well behaved, if a local minimizer is used in conjunction with `nag_glopt_bnd_pso` (e05sa), then it is more likely that the final result will at least be in the near vicinity of a local optimum, and due to the global search characteristics of the particle swarm, this solution should be superior to many other local optima.

Caution should be used in accelerating the rate of convergence, as with faster convergence, less of the domain will remain searchable by the swarm, making it increasingly difficult for the algorithm to detect the basin of attraction of superior local optima. Using the options **Repulsion Initialize** and **Repulsion Finalize** described in Section 12 will help to overcome this, by causing the swarm to diverge away from the current optimum once no more local improvement is likely.

On successful exit with guaranteed success, **ifail** = 0. This may only happen if a **Target Objective Value** is assigned and is reached by the algorithm.

On successful exit without guaranteed success, **ifail** = 1 is returned. This will happen if another finalization criterion is achieved without the detection of an error.

In both cases, the value of **inform** provides further information as to the cause of the exit.

8 Further Comments

The memory used by `nag_glopt_bnd_pso` (e05sa) is relatively static throughout. As such, `nag_glopt_bnd_pso` (e05sa) may be used in problems with high dimension number ($\mathbf{ndim} > 100$) without the concern of computational resource exhaustion, although the probability of successfully locating the global optimum will decrease dramatically with the increase in dimensionality.

Due to the stochastic nature of the algorithm, the result will vary over multiple runs. This is particularly true if arguments and options are chosen to accelerate convergence at the expense of the global search. However, the option **Repeatability** = ON may be set to initialize the internal random number generator using a preset seed, which will result in identical solutions being obtained.

9 Example

This example uses a particle swarm to find the global minimum of the Schwefel function:

$$\underset{\mathbf{x} \in \mathbb{R}^{ndim}}{\text{minimize}} f = \sum_{i=1}^{ndim} x_i \sin(\sqrt{|x_i|})$$

$$x_i \in (-500, 500), \quad \text{for } i = 1, 2, \dots, \mathbf{ndim}.$$

In two dimensions the optimum is $f_{\min} = -837.966$, located at $\mathbf{x} = (-420.97, -420.97)$.

The example demonstrates how to initialize and set the options arrays using `nag_glopt_optset` (e05zk), how to query options using `nag_glopt_optget` (e05zl), and finally how to search for the global optimum using `nag_glopt_bnd_pso` (e05sa). The function is minimized several times to demonstrate using `nag_glopt_bnd_pso` (e05sa) alone, and coupled with local minimizers. This program uses the non-default option **Repeatability** = ON to produce repeatable solutions.

9.1 Program Text

```
function e05sa_example

fprintf('e05sa example results\n\n');

bu = [ 500,  500];
bl = [-500, -500];

npar = nag_int(5);

x_target = [-420.9687463599820;-420.9687463599820];
f_target = -837.9657745448674;

iopts = zeros(100, 1, nag_int_name);
opts = zeros(100, 1);

% Initialize the option arrays for e05sa
[iopts, opts, ifail] = e05zk('Initialize = e05sa', iopts, opts);

% Query some default option values.
[ivalue, rvalue, boundary, optype, ifail] = ...
    e05zl(...
        'Boundary', iopts, opts);
[maxits, rvalue, itsstr, optype, ifail] = ...
    e05zl(...
        'Maximum Iterations Completed', iopts, opts);
[ivalue, distol, cvalue, optype, ifail] = ...
    e05zl(...
        'Distance Tolerance', iopts, opts);
fprintf('\nDefault Option Queries:\n\n');
fprintf('Boundary           : %s\n', boundary);
fprintf('Maximum Iterations Completed : %d (%s)\n', maxits, strtrim(itsstr));
fprintf('Distance Tolerance      : %10.4e\n', distol);

fprintf('\n1. Solution without using coupled local minimizer.\n');
```

```

% Set various options to non-default values if required.
[iopts, opts, ifail] = e05zk(...
    'Repeatability = On', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Verify Gradients = Off', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Boundary = Hyperspherical', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Maximum iterations static = 150', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Repulsion Initialize = 30', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Repulsion Finalize = 30', iopts, opts);

% Catch warnings and assume ifail=1,3 give a reasonable estimate
wstat = warning();
warning('OFF');
% Call e05sa to search for the global optimum.
[xb, fb, iopts, opts, user, itt, inform, ifail] = ...
    e05sa(bl, bu, @objfun, 'e05sxm', iopts, opts, 'npar', npar);

switch ifail
case {0,1}
    % e05sa encountered no errors during operation,
    % and will have returned the best optimum found.
    display_result(x_target, f_target, xb,fb,itt,inform);
case 3
    % An instruction to exit was received by e05sa from objfun or monmod.
    % The exit flag will have been returned in inform.
    display_result(x_target, f_target, xb,fb,itt,inform);
otherwise
    % An error was detected, and a warning has been displayed
end

fprintf('\n2. Solution using coupled local minimizer e04cb.\n');

% Set an objective target
optstr = sprintf('Target Objective Value = %32.16e', f_target);
[iopts, opts, ifail] = e05zk(optstr, iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Target Objective Tolerance = 1.0e-5', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Target Objective Safeguard = 1.0e-8', iopts, opts);

% Set the local minimizer to be e04cb and set corresponding options
[iopts, opts, ifail] = e05zk(...
    'Local Minimizer = e04cb', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Local Interior Iterations = 10', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Local Exterior Iterations = 20', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Local Interior Tolerance = 1.0e-4', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Local Exterior Tolerance = 1.0e-4', iopts, opts);

% Call e05sa to search for the global optimum.
[xb, fb, iopts, opts, user, itt, inform, ifail] = ...
    e05sa(...
        bl, bu, @objfun, 'e05sxm', iopts, opts, 'npar', npar);

switch ifail
case {0,1}
    % e05sa encountered no errors during operation,
    % and will have returned the best optimum found.
    display_result(x_target, f_target, xb,fb,itt,inform);
case 3
    % An instruction to exit was received by e05sa from objfun or monmod.
    % The exit flag will have been returned in inform.
    display_result(x_target, f_target, xb,fb,itt,inform);
otherwise

```

```

    % An error was detected, and a warning has been displayed
end

fprintf('\n3. Solution using coupled local minimizer e04dg.\n');

% Set the local minimizer to be e04dg and set corresponding options
[iopts, opts, ifail] = e05zk(...
    'Local Minimizer = e04dg', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Local Interior Iterations = 5', iopts, opts);
[iopts, opts, ifail] = e05zk(...
    'Local Exterior Iterations = 20', iopts, opts);

%iopts, opts, ifail] = e05zk('Verify Gradients = off', iopts, opts);
% Call e05sa to search for the global optimum.
[xb, fb, iopts, opts, user, itt, inform, ifail] = ...
    e05sa(...
        bl, bu, @objfun, 'e05sxm', iopts, opts, 'npar', npar);

switch ifail
case {0,1}
    % e05sa encountered no errors during operation,
    % and will have returned the best optimum found.
    display_result(x_target, f_target, xb,fb,itt,inform);
case 3
    % An instruction to exit was received by e05sa from objfun or monmod.
    % The exit flag will have been returned in inform.
    display_result(x_target, f_target, xb,fb,itt,inform);
otherwise
    % An error was detected, and a warning has been displayed
end

warning(wstat);

function [mode, objf, vecout, user] = ...
objfun(mode, ndim, x, objf, vecout, nstate, user)
% Test nstate to indicate what stage of computation has been reached.
switch nstate
case (2)
    % objfun is called for the very first time.
case (-1)
    % objfun is called for the first time on a new slave thread.
    % This will never happen if running in a serial context.
case (1)
    % objfun is called on entry to a NAG local minimiser.
case (0)
    % This will be the normal value of NSTATE.
otherwise
    % This is extremely unlikely, and indicates that an error has
    % occurred on the system
    mode = nag_int(-1);
    error('*** Error detected in objfun');
end

% Test mode to determine whether to calculate objf and/or objgrd.
evalobjf = false;
evalobjg = false;
switch mode
case {0,5}
    % Only the value of the objective function is needed.
    evalobjf = true;
case {1,6}
    % Only the values of the NDIM gradients are required.
    evalobjg = true;
case {2,7}
    % Both the objective function and the NDIM gradients are required.
    evalobjf = true;
    evalobjg = true;
otherwise
    mode = nag_int(-1);
    error('*** Illegal value of mode (%d) in objfun', mode);

```

```

end

if evalobjf
    % Evaluate the objective function.
    objf = sum(x(1:double(ndim)).*sin(sqrt(abs(x(1:double(ndim))))));
end

if evalobjg
    % Calculate the gradient of the objective function,
    % and return the result in vecout.
    vecout = sqrt(abs(x));
    for i=1:double(ndim)
        if abs(x(i)) < x02aj
            vecout(i) = 0;
        else
            v = vecout(i);
            vecout(i) = sin(v) + signtransfer(x(i)*cos(v)/(2*v), x(i));
        end
    end
end

function [r] = signtransfer(x, y)
    if y >= 0
        r = abs(x);
    else
        r = -abs(x);
    end
end

function [] = display_result(x_target, f_target, xb,fb,itt,inform)

% Display final counters.
fprintf('\nAlgorithm Statistics\n-----\n');
fprintf('Total complete iterations           : %4d\n', itt(1));
fprintf('Complete iterations since improvement : %4d\n', itt(2));
fprintf('Total particles converged to xb         : %4d\n', itt(3));
fprintf('Total improvements to global optimum   : %4d\n', itt(4));
fprintf('Total function evaluations             : %4d\n', itt(5));
fprintf('Total particles re-initialized         : %4d\n\n', itt(6));

% Display why finalization occurred.
switch inform
case 0
    fprintf('Solution Status : An error was detected by e05sa\n');
case 1
    fprintf('Solution Status : Target value achieved\n');
case 2
    fprintf('Solution Status : Minimum swarm standard deviation obtained\n');
case 3
    fprintf('Solution Status : Sufficient particles converged\n');
case 4
    fprintf('Solution Status : No improvement in preset iteration limit\n');
case 5
    fprintf('Solution Status : Maximum complete iterations attained\n');
case 6
    fprintf('Solution Status : Maximum function evaluations exceeded\n');
otherwise
    fprintf('User termination case:  %d\n', inform);
end

% Display final objective value and location.
fprintf('\n Known objective optimum : %13.5f\n', f_target);
fprintf(' Achieved objective value : %13.5f\n\n', fb);

title = 'Comparison between known and achieved optima.';
clabs = {'x_target'; 'xb          '};
ncols = nag_int(80);
indent = nag_int(0);
[ifail] = x04cb(...
    'G', 'N', [x_target, xb], 'f9.2', title, 'I', clabs, ...
    'C', clabs, ncols, indent);
fprintf('\n');

```

9.2 Program Results

e05sa example results

Default Option Queries:

```
Boundary                : FLOATING
Maximum Iterations Completed : 1000 (DEFAULT)
Distance Tolerance       : 1.0000e-04
```

1. Solution without using coupled local minimizer.

Algorithm Statistics

```
-----
Total complete iterations      : 395
Complete iterations since improvement : 152
Total particles converged to xb : 2
Total improvements to global optimum : 59
Total function evaluations     : 2773
Total particles re-initialized  : 2
```

Solution Status : No improvement in preset iteration limit

```
Known objective optimum : -837.96577
Achieved objective value : -837.96567
```

Comparison between known and achieved optima.

```
  x_target      xb
1  -420.97    -420.95
2  -420.97    -420.94
```

2. Solution using coupled local minimizer e04cb.

Algorithm Statistics

```
-----
Total complete iterations      : 51
Complete iterations since improvement : 1
Total particles converged to xb : 0
Total improvements to global optimum : 12
Total function evaluations     : 537
Total particles re-initialized  : 0
```

Solution Status : Target value achieved

```
Known objective optimum : -837.96577
Achieved objective value : -837.96577
```

Comparison between known and achieved optima.

```
  x_target      xb
1  -420.97    -420.97
2  -420.97    -420.97
```

3. Solution using coupled local minimizer e04dg.

Algorithm Statistics

```
-----
Total complete iterations      : 8
Complete iterations since improvement : 1
Total particles converged to xb : 0
Total improvements to global optimum : 10
Total function evaluations     : 120
Total particles re-initialized  : 0
```

Solution Status : Target value achieved

```
Known objective optimum : -837.96577
Achieved objective value : -837.96561
```

Comparison between known and achieved optima.

	x_target	xb
1	-420.97	-420.94
2	-420.97	-420.98

10 Algorithmic Details

The following pseudo-code describes the algorithm used with the repulsion mechanism.

```

INITIALIZE for  $j = 1, n_p$ 
     $\mathbf{x}_j = \mathbf{R} \in U(\ell_{\text{box}}, \mathbf{u}_{\text{box}})$ 
     $\hat{\mathbf{x}}_j = \mathbf{R} \in U(\ell_{\text{box}}, \mathbf{u}_{\text{box}})$ 
     $\mathbf{v}_j = \mathbf{R} \in U(-\mathbf{V}_{\text{max}}, \mathbf{V}_{\text{max}})$ 
     $\hat{f}_j = F(\hat{\mathbf{x}}_j)$ 
    initialize  $w_j$ 
     $w_j = \begin{cases} W_{\text{max}} & \text{Weight Initialize = MAXIMUM} \\ W_{\text{ini}} & \text{Weight Initialize = INITIAL} \\ R \in U(W_{\text{min}}, W_{\text{max}}) & \text{Weight Initialize = RANDOMIZED} \end{cases}$ 
end for
 $\tilde{\mathbf{x}} = \frac{1}{2}(\ell_{\text{box}} + \mathbf{u}_{\text{box}})$ 
 $\tilde{f} = F(\tilde{\mathbf{x}})$ 
 $I_c = I_s = 0$ 

SWARM while (not finalized),
     $I_c = I_c + 1$ 
    for  $j = 1, n_p$ 
         $\mathbf{x}_j = \text{BOUNDARY}(\mathbf{x}_j, \ell_{\text{box}}, \mathbf{u}_{\text{box}})$ 
         $f_j = F(\mathbf{x}_j)$ 
        if ( $f_j < \hat{f}_j$ )
             $\hat{f}_j = f_j, \hat{\mathbf{x}}_j = \mathbf{x}_j$ 
        if ( $f_j < \tilde{f}$ )
             $\tilde{f} = f_j, \tilde{\mathbf{x}} = \mathbf{x}_j$ 
        end for
        if (new( $\tilde{f}$ ))
            LOCMIN( $\tilde{\mathbf{x}}, \tilde{f}, O_i$ ),  $I_s = 0$ 
            [see note on repulsion below for code insertion]
        else
             $I_s = I_s + 1$ 
        for  $j = 1, n_p$ 
             $\mathbf{v}_j = w_j \mathbf{v}_j + C_s \mathbf{D}_1(\hat{\mathbf{x}}_j - \mathbf{x}_j) + C_g \mathbf{D}_2(\tilde{\mathbf{x}} - \mathbf{x}_j)$ 
             $\mathbf{x}_j = \mathbf{x}_j + \mathbf{v}_j$ 
            if ( $\|\mathbf{x}_j - \tilde{\mathbf{x}}\| < dtol$ )
                reset  $\mathbf{x}_j, \mathbf{v}_j, w_j; \hat{\mathbf{x}}_j = \mathbf{x}_j$ 
            else
                update ( $w_j$ )
            end for
            if (target achieved or termination criterion satisfied)
                finalized = true
            monmod( $\mathbf{x}_j$ )
        end
        LOCMIN( $\tilde{\mathbf{x}}, \tilde{f}, O_e$ )

```

The definition of terms used in the above pseudo-code are as follows.

n_p the number of particles, **npar**
 ℓ_{box} array of **ndim** lower box bounds
 \mathbf{u}_{box} array of **ndim** upper box bounds

\mathbf{x}_j	position of particle j
$\hat{\mathbf{x}}_j$	best position found by particle j
$\tilde{\mathbf{x}}$	best position found by any particle
f_j	$F(\mathbf{x}_j)$
\hat{f}_j	$F(\hat{\mathbf{x}}_j)$, best value found by particle j
\tilde{f}	$F(\tilde{\mathbf{x}})$, best value found by any particle
\mathbf{v}_j	velocity of particle j
w_j	weight on \mathbf{v}_j for velocity update, decreasing according to Weight Decrease
V_{\max}	maximum absolute velocity, dependent upon Maximum Variable Velocity
I_c	swarm iteration counter
I_s	iterations since $\tilde{\mathbf{x}}$ was updated
$\mathbf{D}_1, \mathbf{D}_2$	diagonal matrices with random elements in range $(0, 1)$
C_s	the cognitive advance coefficient which weights velocity towards $\hat{\mathbf{x}}_j$, adjusted using Advance Cognitive
C_g	the global advance coefficient which weights velocity towards $\tilde{\mathbf{x}}$, adjusted using Advance Global
$dtol$	the Distance Tolerance for resetting a converged particle
$\mathbf{R} \in U(\ell_{\text{box}}, \mathbf{u}_{\text{box}})$	an array of random numbers whose i -th element is drawn from a uniform distribution in the range $(\ell_{\text{box}i}, \mathbf{u}_{\text{box}i})$, for $i = 1, 2, \dots, \mathbf{ndim}$
O_i	local optimizer interior options
O_e	local optimizer exterior options
LOCMIN(\mathbf{x}, f, O)	apply local optimizer using the set of options O using the solution (\mathbf{x}, f) as the starting point, if used (not default)
monmod	monitor progress and possibly modify \mathbf{x}_j
BOUNDARY	apply required behaviour for \mathbf{x}_j outside bounding box, (see Boundary)
new(\tilde{f})	true if $\tilde{\mathbf{x}}, \tilde{\mathbf{c}}, \tilde{f}$ were updated at this iteration

Additionally a repulsion phase can be introduced by changing from the default values of options **Repulsion Finalize** (r_f), **Repulsion Initialize** (r_i) and **Repulsion Particles** (r_p). If the number of static particles is denoted n_s then the following can be inserted after the new(\tilde{f}) check in the pseudo-code above.

```

else if ( $n_s \geq r_p$  and  $r_i \leq I_s \leq r_i + r_f$ )
    LOCMIN( $\tilde{\mathbf{x}}, \tilde{f}, O_i$ )
    use  $-C_g$  instead of  $C_g$  in velocity updates
if ( $I_s = r_i + r_f$ )
     $I_s = 0$ 

```

11 Optional Parameters

This section can be skipped if you wish to use the default values for all optional parameters, otherwise, the following is a list of the optional parameters available and a full description of each optional parameter is provided in Section 12.1.

Advance Cognitive
Advance Global
Boundary
Distance Scaling
Distance Tolerance
Function Precision
Local Boundary Restriction
Local Exterior Iterations
Local Exterior Major Iterations
Local Exterior Minor Iterations
Local Exterior Tolerance
Local Interior Iterations
Local Interior Major Iterations
Local Interior Minor Iterations
Local Interior Tolerance
Local Minimizer
Maximum Function Evaluations
Maximum Iterations Completed
Maximum Iterations Static
Maximum Iterations Static Particles
Maximum Particles Converged
Maximum Particles Reset
Maximum Variable Velocity
Optimize
Repeatability
Repulsion Finalize
Repulsion Initialize
Repulsion Particles
Swarm Standard Deviation
Target Objective
Target Objective Safeguard
Target Objective Tolerance
Target Objective Value
Target Warning
Verify Gradients
Weight Decrease
Weight Initial
Weight Initialize
Weight Maximum
Weight Minimum
Weight Reset
Weight Value

11.1 Description of the Optional Parameters

For each option, we give a summary line, a description of the optional parameter and details of constraints.

The summary line contains:

the keywords;

a parameter value, where the letters a , i and r denote options that take character, integer and real values respectively;

the default value, where the symbol ϵ is a generic notation for *machine precision* (see nag_machine_precision (x02aj)), and *Imax* represents the largest representable integer value (see nag_machine_integer_max (x02bb)).

All options accept the value ‘DEFAULT’ in order to return single options to their default states.

Keywords and character values are case insensitive, however they must be separated by at least one space.

For nag_glopt_bnd_pso (e05sa) the maximum length of the argument **cvalue** used by nag_glopt_optget (e05zl) is 15.

Advance Cognitive r Default = 2.0

The cognitive advance coefficient, C_s . When larger than the global advance coefficient, this will cause particles to be attracted toward their previous best positions. Setting $r = 0.0$ will cause nag_glopt_bnd_pso (e05sa) to act predominantly as a local optimizer. Setting $r > 2.0$ may cause the swarm to diverge, and is generally inadvisable. At least one of the global and cognitive coefficients must be nonzero.

Advance Global r Default = 2.0

The global advance coefficient, C_g . When larger than the cognitive coefficient this will encourage convergence toward the best solution yet found. Values $r \in (0, 1)$ will inhibit particles overshooting the optimum. Values $r \in [1, 2)$ cause particles to fly over the optimum some of the time. Larger values can prohibit convergence. Setting $r = 0.0$ will remove any attraction to the current optimum, effectively generating a Monte–Carlo multi-start optimization algorithm. At least one of the global and cognitive coefficients must be nonzero.

Boundary a Default = FLOATING

Determines the behaviour if particles leave the domain described by the box bounds. This only affects the general PSO algorithm, and will not pass down to any NAG local minimizers chosen.

This option is only effective in those dimensions for which $\mathbf{bl}(i) \neq \mathbf{bu}(i)$, $i = 1, 2, \dots, \mathbf{ndim}$.

IGNORE

The box bounds are ignored. The objective function is still evaluated at the new particle position.

RESET

The particle is re-initialized inside the domain. $\hat{\mathbf{x}}_j$ and \hat{f}_j are not affected.

FLOATING

The particle position remains the same, however the objective function will not be evaluated at the next iteration. The particle will probably be advected back into the domain at the next advance due to attraction by the cognitive and global memory.

HYPERSPHERICAL

The box bounds are wrapped around an $ndim$ -dimensional hypersphere. As such a particle leaving through a lower bound will immediately re-enter through the corresponding upper bound and vice versa. The standard distance between particles is also modified accordingly.

FIXED

The particle rests on the boundary, with the corresponding dimensional velocity set to 0.0.

Distance Scaling*a*

Default = ON

Determines whether distances should be scaled by box widths.

ON

When a distance is calculated between \mathbf{x} and \mathbf{y} , a scaled L^2 norm is used.

$$L^2(\mathbf{x}, \mathbf{y}) = \left(\sum_{\{i | \mathbf{u}_i \neq \mathbf{l}_i, i \leq \text{ndim}\}} \left(\frac{x_i - y_i}{\mathbf{u}_i - \mathbf{l}_i} \right)^2 \right)^{\frac{1}{2}}.$$

OFF

Distances are calculated as the standard L^2 norm without any rescaling.

$$L^2(\mathbf{x}, \mathbf{y}) = \left(\sum_{i=1}^{\text{ndim}} (x_i - y_i)^2 \right)^{\frac{1}{2}}.$$

Distance Tolerance*r*Default = 10^{-4}

This is the distance, *dtol* between particles and the global optimum which must be reached for the particle to be considered converged, i.e., that any subsequent movement of such a particle cannot significantly alter the global optimum. Once achieved the particle is reset into the box bounds to continue searching.

Constraint: $r > 0.0$.

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

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

The value of ϵ_r should reflect the relative precision of $1 + |F(\mathbf{x})|$; i.e., ϵ_r acts as a relative precision when $|F|$ is large, and as an absolute precision when $|F|$ is small. For example, if $F(\mathbf{x})$ is typically of order 1000 and the first six significant digits are known to be correct, an appropriate value for ϵ_r would be 10^{-6} . In contrast, if $F(\mathbf{x})$ is typically of order 10^{-4} and the first six significant digits are known to be correct, an appropriate value for ϵ_r would be 10^{-10} . The choice of ϵ_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 ϵ_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.

Local Boundary Restriction*r*

Default = 0.5

Contracts the box boundaries used by a box constrained local minimizer to, $[\beta_l, \beta_u]$, containing the start point x , where

$$\begin{aligned} \partial_i &= r \times (\mathbf{u}_i - \mathbf{l}_i) \\ \beta_l^i &= \max(\mathbf{l}_i, x_i - \frac{\partial_i}{2}) \\ \beta_u^i &= \min(\mathbf{u}_i, x_i + \frac{\partial_i}{2}), \quad i = 1, \dots, \text{ndim}. \end{aligned}$$

Smaller values of r thereby restrict the size of the domain exposed to the local minimizer, possibly reducing the amount of work done by the local minimizer.

Constraint: $0.0 \leq r \leq 1.0$.

Local Interior Iterations*i*₁**Local Interior Major Iterations***i*₁**Local Exterior Iterations***i*₂**Local Exterior Major Iterations***i*₂

The maximum number of iterations or function evaluations the chosen local minimizer will perform inside (outside) the main loop if applicable. For the NAG minimizers these correspond to:

Minimizer	Parameter/option	Default Interior	Default Exterior
e04cb	maxcal	ndim + 10	2 × ndim + 15
e04dg	Iteration Limit	max(30, 3 × ndim)	max(50, 5 × ndim)
e04uc	Major Iteration Limit	max(10, 2 × ndim)	max(30, 3 × ndim)

Unless set, these are functions of the parameters passed to nag_glopt_bnd_pso (e05sa).

Setting $i = 0$ will disable the local minimizer in the corresponding algorithmic region. For example, setting **Local Interior Iterations** = 0 and **Local Exterior Iterations** = 30 will cause the algorithm to perform no local minimizations inside the main loop of the algorithm, and a local minimization with upto 30 iterations after the main loop has been exited.

Note: currently nag_opt_bounds_quasi_func_easy (e04jy) or nag_opt_bounds_mod_deriv_easy (e04kz) are restricted to using $400 \times \mathbf{ndim}$ and $50 \times \mathbf{ndim}$ as function evaluation limits respectively. This applies to both local minimizations inside and outside the main loop. They may still be deactivated in either phase by setting $i = 0$, and may subsequently be reactivated in either phase by setting $i > 0$.

Constraint: $i_1 \geq 0, i_2 \geq 0$.

Local Interior Tolerance	r_1	Default = 10^{-4}
Local Exterior Tolerance	r_2	Default = 10^{-4}

This is the tolerance provided to a local minimizer in the interior (exterior) of the main loop of the algorithm.

Constraint: $r_1 > 0.0, r_2 > 0.0$.

Local Interior Minor Iterations	i_1
Local Exterior Minor Iterations	i_2

Where applicable, the secondary number of iterations the chosen local minimizer will use inside (outside) the main loop. Currently the relevant default values are:

Minimizer	Parameter/option	Default Interior	Default Exterior
e04uc	Minor Iteration Limit	max(10, 2 × ndim)	max(30, 3 × ndim)

Constraint: $i_1 \geq 0, i_2 \geq 0$.

Local Minimizer	a	Default = OFF
------------------------	-----	---------------

Allows for a choice of Chapter E04 functions to be used as a coupled, dedicated local minimizer.

OFF

No local minimization will be performed in either the INTERIOR or EXTERIOR sections of the algorithm.

nag_opt_uncon_simplex (e04cb)

Use nag_opt_uncon_simplex (e04cb) as the local minimizer. This does not require the calculation of derivatives.

On a call to **objfun** during a local minimization, **mode** = 5.

nag_opt_bounds_mod_deriv_easy (e04kz)

Use nag_opt_bounds_mod_deriv_easy (e04kz) as the local minimizer. This requires the calculation of derivatives in **objfun**, as indicated by **mode**.

The box bounds forwarded to this function from nag_glopt_bnd_pso (e05sa) will have been acted upon by **Local Boundary Restriction**. As such, the domain exposed may be greatly smaller than that provided to nag_glopt_bnd_pso (e05sa).

Accurate derivatives must be provided to this function, and will not be approximated internally. Each iteration of this local minimizer also requires the calculation of both the objective function and its derivative. Hence on a call to **objfun** during a local minimization, **mode** = 7.

nag_opt_bounds_quasi_func_easy (e04jy)

Use nag_opt_bounds_quasi_func_easy (e04jy) as the local minimizer. This does not require the calculation of derivatives.

On a call to **objfun** during a local minimization, **mode** = 5.

The box bounds forwarded to this function from nag_glopt_bnd_pso (e05sa) will have been acted upon by **Local Boundary Restriction**. As such, the domain exposed may be greatly smaller than that provided to nag_glopt_bnd_pso (e05sa).

nag_opt_uncon_conjgrd_comp (e04dg)

nag_opt_uncon_conjgrd_comp (e04dg)

Use nag_opt_uncon_conjgrd_comp (e04dg) as the local minimizer.

Accurate derivatives must be provided, and will not be approximated internally. Additionally, each call to **objfun** during a local minimization will require either the objective to be evaluated alone, or both the objective and its gradient to be evaluated. Hence on a call to **objfun**, **mode** = 5 or 7.

nag_opt_nlp1_solve (e04uc)

nag_opt_nlp1_solve (e04uc)

Use nag_opt_nlp1_solve (e04uc) as the local minimizer. This operates such that any derivatives of the objective function that you cannot supply, will be approximated internally using finite differences.

Either, the objective, objective gradient, or both may be requested during a local minimization, and as such on a call to **objfun**, **mode** = 1, 2 or 5.

The box bounds forwarded to this function from nag_glopt_bnd_pso (e05sa) will have been acted upon by **Local Boundary Restriction**. As such, the domain exposed may be greatly smaller than that provided to nag_glopt_bnd_pso (e05sa).

Maximum Function Evaluations

i

Default = *Imax*

The maximum number of evaluations of the objective function. When reached this will return **ifail** = 1 and **inform** = 6.

Constraint: i > 0.

Maximum Iterations Completed

i

Default = 1000 × **ndim**

The maximum number of complete iterations that may be performed. Once exceeded nag_glopt_bnd_pso (e05sa) will exit with **ifail** = 1 and **inform** = 5.

Unless set, this adapts to the parameters passed to nag_glopt_bnd_pso (e05sa).

Constraint: i ≥ 1.

Maximum Iterations Static

i

Default = 100

The maximum number of iterations without any improvement to the current global optimum. If exceeded nag_glopt_bnd_pso (e05sa) will exit with **ifail** = 1 and **inform** = 4. This exit will be hindered by setting **Maximum Iterations Static Particles** to larger values.

Constraint: i ≥ 1.

Maximum Iterations Static Particles

i

Default = 0

The minimum number of particles that must have converged to the current optimum before the function may exit due to **Maximum Iterations Static** with **ifail** = 1 and **inform** = 4.

Constraint: i ≥ 0.

Maximum Particles Converged *i* Default = *Imax*

The maximum number of particles that may converge to the current optimum. When achieved, `nag_glopt_bnd_pso` (e05sa) will exit with **ifail** = 1 and **inform** = 3. This exit will be hindered by setting ‘**Repulsion**’ options, as these cause the swarm to re-expand.

Constraint: $i > 0$.

Maximum Particles Reset *i* Default = *Imax*

The maximum number of particles that may be reset after converging to the current optimum. Once achieved no further particles will be reset, and any particles within **Distance Tolerance** of the global optimum will continue to evolve as normal.

Constraint: $i > 0$.

Maximum Variable Velocity *r* Default = 0.25

Along any dimension j , the absolute velocity is bounded above by $|\mathbf{v}_j| \leq r \times (\mathbf{u}_j - \ell_j) = \mathbf{V}_{\max}$. Very low values will greatly increase convergence time. There is no upper limit, although larger values will allow more particles to be advected out of the box bounds, and values greater than 4.0 may cause significant and potentially unrecoverable swarm divergence.

Constraint: $r > 0.0$.

Optimize *a* Default = MINIMIZE

Determines whether to maximize or minimize the objective function.

MINIMIZE

The objective function will be minimized.

MAXIMIZE

The objective function will be maximized. This is accomplished by minimizing the negative of the objective.

Repeatability *a* Default = OFF

Allows for the same random number generator seed to be used for every call to `nag_glopt_bnd_pso` (e05sa). **Repeatability** = OFF is recommended in general.

OFF

The internal generation of random numbers will be nonrepeatable.

ON

The same seed will be used.

Repulsion Finalize *i* Default = *Imax*

The number of iterations performed in a repulsive phase before re-contraction. This allows a re-diversified swarm to contract back toward the current optimum, allowing for a finer search of the near optimum space.

Constraint: $i \geq 2$.

Repulsion Initialize *i* Default = *Imax*

The number of iterations without any improvement to the global optimum before the algorithm begins a repulsive phase. This phase allows the particle swarm to re-expand away from the current optimum, allowing more of the domain to be investigated. The repulsive phase is automatically ended if a superior optimum is found.

Constraint: $i \geq 2$.

Repulsion Particles i Default = 0

The number of particles required to have converged to the current optimum before any repulsive phase may be initialized. This will prevent repulsion before a satisfactory search of the near optimum area has been performed, which may happen for large dimensional problems.

Constraint: $i \geq 0$.

Swarm Standard Deviation r Default = 0.1

The target standard deviation of the particle distances from the current optimum. Once the standard deviation is below this level, `nag_glopt_bnd_pso` (e05sa) will exit with **ifail** = 1 and **inform** = 2. This criterion will be penalized by the use of ‘**Repulsion**’ options, as these cause the swarm to re-expand, increasing the standard deviation of the particle distances from the best point.

In SMP parallel implementations of `nag_glopt_bnd_pso` (e05sa), the standard deviation will be calculated based only on the particles local to the particular thread that checks for finalization. Considerably fewer particles may be used in this calculation than when the algorithm is run in serial. It is therefore recommended that you provide a smaller value of **Swarm Standard Deviation** when running in parallel than when running in serial.

Constraint: $r \geq 0.0$.

Target Objective a Default = OFF
Target Objective Value r Default = 0.0

Activate or deactivate the use of a target value as a finalization criterion. If active, then once the supplied target value for the objective function is found (beyond the first iteration if **Target Warning** is active) `nag_glopt_bnd_pso` (e05sa) will exit with **ifail** = 0 and **inform** = 1. Other than checking for feasibility only (**Optimize** = CONSTRAINTS), this is the only finalization criterion that guarantees that the algorithm has been successful. If the target value was achieved at the initialization phase or first iteration and **Target Warning** is active, `nag_glopt_bnd_pso` (e05sa) will exit with **ifail** = 2. This option may take any real value r , or the character ON/OFF as well as DEFAULT. If this option is queried using `nag_glopt_optget` (e05zl), the current value of r will be returned in **rvalue**, and **cvalue** will indicate whether this option is ON or OFF. The behaviour of the option is as follows:

r

Once a point is found with an objective value within the **Target Objective Tolerance** of r , `nag_glopt_bnd_pso` (e05sa) will exit successfully with **ifail** = 0 and **inform** = 1.

OFF

The current value of r will remain stored, however it will not be used as a finalization criterion.

ON

The current value of r stored will be used as a finalization criterion.

DEFAULT

The stored value of r will be reset to its default value (0.0), and this finalization criterion will be deactivated.

Target Objective Safeguard r Default = 100.0e

If you have given a target objective value to reach in *objval* (the value of the optional parameter **Target Objective Value**), *objsf* sets your desired safeguarded termination tolerance, for when *objval* is close to zero.

Constraint: $objsf \geq 2\epsilon$.

Target Objective Tolerance r Default = 0.0

The optional tolerance to a user-specified target value.

Constraint: $r \geq 0.0$.

Target Warning *a* Default = OFF

Activates or deactivates the error exit associated with the target value being achieved before entry into the main loop of the algorithm, **ifail** = 2.

OFF

No error will be returned, and the function will exit normally.

ON

An error will be returned if the target objective is reached prematurely, and the function will exit with **ifail** = 2.

Verify Gradients *a* Default = ON

Adjusts the level of gradient checking performed when gradients are required. Gradient checks are only performed on the first call to the chosen local minimizer if it requires gradients. There is no guarantee that the gradient check will be correct, as the finite differences used in the gradient check are themselves subject to inaccuracies.

OFF

No gradient checking will be performed.

ON

A cheap gradient check will be performed on both the gradients corresponding to the objective through **objfun**.

OBJECTIVE

FULL

A more expensive gradient check will be performed on the gradients corresponding to the objective **objfun**.

Weight Decrease *a* Default = INTEREST

Determines how particle weights decrease.

OFF

Weights do not decrease.

INTEREST

Weights decrease through compound interest as $w_{IT+1} = w_{IT}(1 - W_{val})$, where W_{val} is the **Weight Value** and IT is the current number of iterations.

LINEAR

Weights decrease linearly following $w_{IT+1} = w_{IT} - IT \times (W_{max} - W_{min})/IT_{max}$, where IT is the iteration number and IT_{max} is the maximum number of iterations as set by **Maximum Iterations Completed**.

Weight Initial *r* Default = W_{max}

The initial value of any particle's inertial weight, W_{ini} , or the minimum possible initial value if initial weights are randomized. When set, this will override **Weight Initialize** = RANDOMIZED or MAXIMUM, and as such these must be set afterwards if so desired.

Constraint: $W_{min} \leq r \leq W_{max}$.

Weight Initialize *a* Default = MAXIMUM

Determines how the initial weights are distributed.

INITIAL

All weights are initialized at the initial weight, W_{ini} , if set. If **Weight Initial** has not been set, this will be the maximum weight, W_{max} .

MAXIMUM

All weights are initialized at the maximum weight, W_{max} .

RANDOMIZED

Weights are uniformly distributed in (W_{min}, W_{max}) or (W_{ini}, W_{max}) if **Weight Initial** has been set.

Weight Maximum r

Default = 1.0

The maximum particle weight, W_{max} .

Constraint: $1.0 \geq r \geq W_{min}$ (If W_{ini} has been set then $1.0 \geq r \geq W_{ini}$.)

Weight Minimum r

Default = 0.1

The minimum achievable weight of any particle, W_{min} . Once achieved, no further weight reduction is possible.

Constraint: $0.0 \leq r \leq W_{max}$ (If W_{ini} has been set then $0.0 \leq r \leq W_{ini}$.)

Weight Reset a

Default = MAXIMUM

Determines how particle weights are re-initialized.

INITIAL

Weights are re-initialized at the initial weight if set. If **Weight Initial** has not been set, this will be the maximum weight.

MAXIMUM

Weights are re-initialized at the maximum weight.

RANDOMIZED

Weights are uniformly distributed in (W_{min}, W_{max}) or (W_{ini}, W_{max}) if **Weight Initial** has been set.

Weight Value r

Default = 0.01

The constant W_{val} used with **Weight Decrease** = INTEREST.

Constraint: $0.0 \leq r \leq \frac{1}{3}$.
