


# Genetic Algorithm and Direct Search Toolbox 2

## User's Guide



**MATLAB<sup>®</sup>**

## How to Contact The MathWorks



www.mathworks.com  
comp.soft-sys.matlab  
www.mathworks.com/contact\_TS.html

Web  
Newsgroup  
Technical Support



suggest@mathworks.com  
bugs@mathworks.com  
doc@mathworks.com  
service@mathworks.com  
info@mathworks.com

Product enhancement suggestions  
Bug reports  
Documentation error reports  
Order status, license renewals, passcodes  
Sales, pricing, and general information



508-647-7000 (Phone)



508-647-7001 (Fax)



The MathWorks, Inc.  
3 Apple Hill Drive  
Natick, MA 01760-2098

For contact information about worldwide offices, see the MathWorks Web site.

*Genetic Algorithm and Direct Search Toolbox User's Guide*

© COPYRIGHT 2004–2007 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

**FEDERAL ACQUISITION:** This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

### Trademarks

MATLAB, Simulink, Stateflow, Handle Graphics, Real-Time Workshop, and xPC TargetBox are registered trademarks, and SimBiology, SimEvents, and SimHydraulics are trademarks of The MathWorks, Inc.

Other product or brand names are trademarks or registered trademarks of their respective holders.

### Patents

The MathWorks products are protected by one or more U.S. patents. Please see [www.mathworks.com/patents](http://www.mathworks.com/patents) for more information.

**Revision History**

|                |                 |   |
|----------------|-----------------|---|
| January 2004   | Online only     | New for Version 1.0 (Release 13SP1+)      |
| June 2004      | First printing  | Revised for Version 1.0.1 (Release 14)    |
| October 2004   | Online only     | Revised for Version 1.0.2 (Release 14SP1) |
| March 2005     | Online only     | Revised for Version 1.0.3 (Release 14SP2) |
| September 2005 | Second printing | Revised for Version 2.0 (Release 14SP3)   |
| March 2006     | Online only     | Revised for Version 2.0.1 (Release 2006a) |
| September 2006 | Online only     | Revised for Version 2.0.2 (Release 2006b) |
| March 2007     | Online only     | Revised for Version 2.1 (Release 2007a)   |



## Acknowledgment

The MathWorks would like to acknowledge Mark A. Abramson for his contributions to the Genetic Algorithm and Direct Search Toolbox. He researched and helped with the development of the linearly constrained pattern search algorithm.

Dr. Mark A. Abramson is the Deputy Head and Assistant Professor in the Department of Mathematics and Statistics, Air Force Institute of Technology, Wright-Patterson AFB, Ohio. Dr. Abramson is actively engaged in research for pattern search methods. He has published over 26 technical papers and has worked on NOMAD, a software package for pattern search.

## Acknowledgment

---

## What Is Genetic Algorithm and Direct Search Toolbox?

**1**

## Writing M-Files for Functions You Want to Optimize

**2**

|   |            |
|---|------------|
| <b>Example — Writing an M-File</b> .....  | <b>2-2</b> |
| <b>Maximizing Versus Minimizing</b> ..... | <b>2-3</b> |

## Getting Started with Direct Search

**3**

|   |             |
|---|-------------|
| <b>What Is Direct Search?</b> .....   | <b>3-2</b>  |
| <b>Performing a Pattern Search</b> .....  | <b>3-3</b>  |
| Calling patternsearch at the Command Line .....                                     | <b>3-3</b>  |
| Using the Pattern Search Tool .....   | <b>3-3</b>  |
| <b>Example: Finding the Minimum of a Function Using<br/>the GPS Algorithm</b> ..... | <b>3-6</b>  |
| Objective Function .....  | <b>3-6</b>  |
| Finding the Minimum of the Function .....   | <b>3-7</b>  |
| Plotting the Objective Function Values and Mesh Sizes ..                            | <b>3-8</b>  |
| <b>Pattern Search Terminology</b> .....   | <b>3-10</b> |
| Patterns .....  | <b>3-10</b> |
| Meshes .....  | <b>3-11</b> |

|   |             |
|---|-------------|
| Polling .....   | 3-12        |
| <b>How Pattern Search Works .....</b>                       | <b>3-14</b> |
| Successful Polls .....                                      | 3-14        |
| An Unsuccessful Poll .....                                  | 3-17        |
| Displaying the Results at Each Iteration .....              | 3-19        |
| More Iterations .....                                       | 3-20        |
| Stopping Conditions for the Pattern Search .....            | 3-21        |
| <b>Description of the Nonlinear Constraint Solver .....</b> | <b>3-23</b> |

## Getting Started with the Genetic Algorithm

# 4

|   |             |
|---|-------------|
| <b>What Is the Genetic Algorithm? .....</b>       | <b>4-2</b>  |
| <b>Using the Genetic Algorithm .....</b>          | <b>4-3</b>  |
| Calling the Function ga at the Command Line ..... | 4-3         |
| Using the Genetic Algorithm Tool .....            | 4-4         |
| <b>Example: Rastrigin's Function .....</b>        | <b>4-7</b>  |
| Rastrigin's Function .....                        | 4-7         |
| Finding the Minimum of Rastrigin's Function ..... | 4-9         |
| Finding the Minimum from the Command Line .....   | 4-11        |
| Displaying Plots .....                            | 4-12        |
| <b>Some Genetic Algorithm Terminology .....</b>   | <b>4-15</b> |
| Fitness Functions .....                           | 4-15        |
| Individuals .....                                 | 4-15        |
| Populations and Generations .....                 | 4-16        |
| Diversity .....                                   | 4-16        |
| Fitness Values and Best Fitness Values .....      | 4-16        |
| Parents and Children .....                        | 4-17        |
| <b>How the Genetic Algorithm Works .....</b>      | <b>4-18</b> |
| Outline of the Algorithm .....                    | 4-18        |
| Initial Population .....                          | 4-19        |
| Creating the Next Generation .....                | 4-20        |



|   |      |
|---|------|
| Plots of Later Generations .....            | 4-22 |
| Stopping Conditions for the Algorithm ..... | 4-22 |

|   |             |
|---|-------------|
| <b>Description of the Nonlinear Constraint Solver .....</b> | <b>4-25</b> |
|---|-------------|

## **Getting Started with Simulated Annealing and Threshold Acceptance**

# **5**

|   |            |
|---|------------|
| <b>What Are Simulated Annealing and Threshold Acceptance? .....</b> | <b>5-2</b> |
|---|------------|

|   |            |
|---|------------|
| <b>Calling the Simulated Annealing and Threshold Acceptance Solvers .....</b> | <b>5-3</b> |
|---|------------|

|   |            |
|---|------------|
| <b>Example: Minimizing De Jong's Fifth Function .....</b> | <b>5-4</b> |
|---|------------|

|  |            |
|--|------------|
| <b>Some Simulated Annealing and Threshold Acceptance</b> |            |
| <b>Terminology .....</b>                                 | <b>5-6</b> |
| Objective Function .....                                 | 5-6        |
| Temperature .....  | 5-6        |
| Annealing Schedule .....                                 | 5-6        |
| Reannealing .....  | 5-6        |

|  |            |
|--|------------|
| <b>How Simulated Annealing and Threshold Acceptance Work .....</b> | <b>5-8</b> |
| Outline of the Algorithm .....                                     | 5-8        |
| Stopping Conditions for the Algorithm .....                        | 5-8        |

## **Using Direct Search**

# **6**

|   |            |
|---|------------|
| <b>Overview of the Pattern Search Tool .....</b>    | <b>6-2</b> |
| Opening the Pattern Search Tool .....               | 6-3        |
| Defining a Problem in the Pattern Search Tool ..... | 6-4        |

|   |      |
|---|------|
| Running a Pattern Search .....                      | 6-5  |
| Example — A Linearly Constrained Problem .....      | 6-6  |
| Pausing and Stopping the Algorithm .....            | 6-9  |
| Displaying Plots .....                              | 6-9  |
| Example — Working with a Custom Plot Function ..... | 6-10 |
| Setting Options in the Pattern Search Tool .....    | 6-15 |
| Importing and Exporting Options and Problems .....  | 6-16 |
| Generating an M-File .....                          | 6-19 |

### **Performing a Pattern Search from the Command**

|  |      |
|--|------|
| <b>Line</b> .....  | 6-20 |
| Calling patternsearch with the Default Options .....             | 6-20 |
| Setting Options for patternsearch at the Command Line ..         | 6-22 |
| Using Options and Problems from the Pattern Search<br>Tool ..... | 6-24 |

### **Pattern Search Examples** .....

|  |      |
|--|------|
| Poll Method .....                                  | 6-26 |
| Complete Poll .....                                | 6-28 |
| Using a Search Method .....                        | 6-32 |
| Mesh Expansion and Contraction .....               | 6-35 |
| Mesh Accelerator .....                             | 6-40 |
| Using Cache .....                                  | 6-41 |
| Setting Tolerances for the Solver .....            | 6-44 |
| Constrained Minimization Using patternsearch ..... | 6-50 |

### **Parameterizing Functions Called by patternsearch** ...

|  |      |
|--|------|
| Parameterizing Functions Using Anonymous Functions<br>with patternsearch ..... | 6-53 |
| Parameterizing a Function Using a Nested Function with<br>patternsearch .....  | 6-55 |

## **Using the Genetic Algorithm**

# **7**

|  |     |
|--|-----|
| <b>Overview of the Genetic Algorithm Tool</b> .....                      | 7-2 |
| Opening the Genetic Algorithm Tool .....                                 | 7-2 |
| Defining an Unconstrained Problem in the Genetic<br>Algorithm Tool ..... | 7-4 |

|   |      |
|---|------|
| Running the Genetic Algorithm .....   | 7-4  |
| Pausing and Stopping the Algorithm .....                                    | 7-6  |
| Displaying Plots .....  | 7-8  |
| Example — Creating a Custom Plot Function .....                             | 7-9  |
| Reproducing Your Results .....  | 7-12 |
| Setting Options in the Genetic Algorithm Tool .....                         | 7-14 |
| Importing and Exporting Options and Problems .....                          | 7-15 |
| Example — Resuming the Genetic Algorithm from the<br>Final Population ..... | 7-19 |
| Generating an M-File .....  | 7-23 |
| <br>  |      |
| <b>Using the Genetic Algorithm from the Command</b>                         |      |
| <b>Line</b> .....   | 7-24 |
| Running ga with the Default Options .....                                   | 7-24 |
| Setting Options for ga at the Command Line .....                            | 7-26 |
| Using Options and Problems from the Genetic Algorithm<br>Tool .....         | 7-28 |
| Reproducing Your Results .....  | 7-29 |
| Resuming ga from the Final Population of a Previous<br>Run .....            | 7-30 |
| Running ga from an M-File .....   | 7-31 |
| <br>  |      |
| <b>Genetic Algorithm Examples</b> .....                                     | 7-34 |
| Population Diversity .....  | 7-34 |
| Fitness Scaling .....   | 7-39 |
| Selection .....   | 7-42 |
| Reproduction Options .....  | 7-43 |
| Mutation and Crossover .....  | 7-43 |
| Setting the Amount of Mutation .....  | 7-44 |
| Setting the Crossover Fraction .....  | 7-46 |
| Comparing Results for Varying Crossover Fractions .....                     | 7-50 |
| Global Versus Local Minima .....  | 7-52 |
| Using a Hybrid Function .....   | 7-56 |
| Setting the Maximum Number of Generations .....                             | 7-59 |
| Vectorizing the Fitness Function .....                                      | 7-60 |
| Constrained Minimization Using ga .....                                     | 7-61 |
| <br>  |      |
| <b>Parameterizing Functions Called by ga</b> .....                          | 7-67 |
| Parameterizing Functions Using Anonymous Functions<br>with ga .....         | 7-67 |
| Parameterizing a Function Using a Nested Function with<br>ga .....          | 7-68 |

## Using the Simulated Annealing and Threshold Acceptance Algorithms

### 8

|  |     |
|--|-----|
| <b>Using the Simulated Annealing and Threshold Acceptance Algorithms from the Command Line</b> . . . | 8-2 |
| Running simulannealbnd and threshacceptbnd with the Default Options . . . . .                        | 8-2 |
| Setting Options for simulannealbnd and threshacceptbnd at the Command Line . . . . .                 | 8-3 |
| Reproducing Your Results . . . . .   | 8-5 |
| <br>   |     |
| <b>Simulated Annealing and Threshold Acceptance Examples</b> . . . . .                               | 8-7 |

## Options Reference

### 9

|   |      |
|---|------|
| <b>Pattern Search Options</b> . . . . .               | 9-2  |
| Plot Options . . . . .                                | 9-3  |
| Poll Options . . . . .                                | 9-5  |
| Search Options . . . . .                              | 9-8  |
| Mesh Options . . . . .                                | 9-12 |
| Algorithm Settings . . . . .                          | 9-13 |
| Cache Options . . . . .                               | 9-14 |
| Stopping Criteria . . . . .                           | 9-14 |
| Output Function Options . . . . .                     | 9-16 |
| Display to Command Window Options . . . . .           | 9-18 |
| Vectorize Option . . . . .                            | 9-18 |
| Options Table for Pattern Search Algorithms . . . . . | 9-19 |
| <br>  |      |
| <b>Genetic Algorithm Options</b> . . . . .            | 9-23 |
| Plot Options . . . . .                                | 9-24 |
| Population Options . . . . .                          | 9-27 |
| Fitness Scaling Options . . . . .                     | 9-29 |
| Selection Options . . . . .                           | 9-31 |
| Reproduction Options . . . . .                        | 9-32 |
| Mutation Options . . . . .                            | 9-33 |
| Crossover Options . . . . .                           | 9-35 |

|   |      |
|---|------|
| Migration Options .....                 | 9-39 |
| Algorithm Settings .....                | 9-40 |
| Hybrid Function Options .....           | 9-40 |
| Stopping Criteria Options .....         | 9-40 |
| Output Function Options .....           | 9-41 |
| Display to Command Window Options ..... | 9-43 |
| Vectorize Option .....                  | 9-44 |

**Simulated Annealing and Threshold Acceptance**

|                                 |      |
|---------------------------------|------|
| <b>Algorithm Options</b> .....  | 9-45 |
| Plot Options .....              | 9-45 |
| Temperature Options .....       | 9-48 |
| Algorithm Settings .....        | 9-49 |
| Hybrid Function Options .....   | 9-49 |
| Stopping Criteria Options ..... | 9-50 |
| Output Function Options .....   | 9-50 |
| Display Options .....           | 9-52 |

**Functions — By Category**

**10**

|  |      |
|--|------|
| <b>Genetic Algorithm</b> .....                                       | 10-2 |
| <b>Direct Search</b> .....   | 10-2 |
| <b>Simulated Annealing and Threshold Acceptance Algorithms</b> ..... | 10-2 |

**Functions — Alphabetical List**

**11**

**Index**



# What Is Genetic Algorithm and Direct Search Toolbox?

---

Genetic Algorithm and Direct Search Toolbox is a collection of functions that extend the capabilities of Optimization Toolbox and the MATLAB® numeric computing environment. Genetic Algorithm and Direct Search Toolbox includes routines for solving optimization problems using

- Direct search
- Genetic algorithm
- Simulated annealing

These algorithms enable you to solve a variety of optimization problems that lie outside the scope of Optimization Toolbox.

All the toolbox functions are MATLAB M-files made up of MATLAB statements that implement specialized optimization algorithms. You can view the MATLAB code for these functions using the statement

```
type function_name
```

You can extend the capabilities of Genetic Algorithm and Direct Search Toolbox by writing your own M-files, or by using the toolbox in combination with other toolboxes, or with MATLAB or Simulink®.





# Writing M-Files for Functions You Want to Optimize

---

To use Genetic Algorithm and Direct Search Toolbox, you must first write an M-file that computes the function you want to optimize. The M-file should accept a vector, whose length is the number of independent variables for the objective function, and return a scalar. This section explains how to write the M-file and covers the following topics:

- “Example — Writing an M-File” on page 2-2
- “Maximizing Versus Minimizing” on page 2-3

## Example – Writing an M-File

The following example shows how to write an M-file for the function you want to optimize. Suppose that you want to minimize the function

$$f(x_1, x_2) = x_1^2 - 2x_1x_2 + 6x_1 + x_2^2 - 6x_2$$

The M-file that computes this function must accept a vector  $x$  of length 2, corresponding to the variables  $x_1$  and  $x_2$ , and return a scalar equal to the value of the function at  $x$ . To write the M-file, do the following steps:

- 1** Select **New** from the **MATLAB File** menu.
- 2** Select **M-File**. This opens a new M-file in the editor.
- 3** In the M-file, enter the following two lines of code:

```
function z = my_fun(x)
z = x(1)^2 - 2*x(1)*x(2) + 6*x(1) + x(2)^2 - 6*x(2);
```

- 4** Save the M-file in a directory on the **MATLAB** path.

To check that the M-file returns the correct value, enter

```
my_fun([2 3])
```

```
ans =
```

```
-5
```

## Maximizing Versus Minimizing

The optimization functions in Genetic Algorithm and Direct Search Toolbox minimize the objective or fitness function. That is, they solve problems of the form

$$\underset{x}{\text{minimize}} \quad f(x)$$

If you want to maximize  $f(x)$ , you can do so by minimizing  $-f(x)$ , because the point at which the minimum of  $-f(x)$  occurs is the same as the point at which the maximum of  $f(x)$  occurs.

For example, suppose you want to maximize the function

$$f(x_1, x_2) = x_1^2 - 2x_1x_2 + 6x_1 + x_2^2 - 6x_2$$

described in the preceding section. In this case, you should write your M-file to compute

$$-f(x_1, x_2) = -x_1^2 + 2x_1x_2 - 6x_1 - x_2^2 + 6x_2$$

and minimize this function.



# Getting Started with Direct Search

---

|   |  |
|---|--|
| What Is Direct Search? (p. 3-2)   | Introduces direct search and pattern search.                                 |
| Performing a Pattern Search (p. 3-3)  | Explains the main function in the toolbox for performing pattern search.     |
| Example: Finding the Minimum of a Function Using the GPS Algorithm (p. 3-6) | Provides an example of solving an optimization problem using pattern search. |
| Pattern Search Terminology (p. 3-10)  | Explains some basic pattern search terminology.                              |
| How Pattern Search Works (p. 3-14)  | Provides an overview of direct search algorithms.                            |
| Description of the Nonlinear Constraint Solver (p. 3-23)                    | Explains the Augmented Lagrangian Pattern Search (ALPS).                     |

## What Is Direct Search?

Direct search is a method for solving optimization problems that does not require any information about the gradient of the objective function. As opposed to more traditional optimization methods that use information about the gradient or higher derivatives to search for an optimal point, a direct search algorithm searches a set of points around the current point, looking for one where the value of the objective function is lower than the value at the current point. You can use direct search to solve problems for which the objective function is not differentiable, stochastic, or even continuous.

The Genetic Algorithm and Direct Search Toolbox implements two direct search algorithms called the generalized pattern search (GPS) algorithm and the mesh adaptive search (MADS) algorithm. Both are *pattern search* algorithms that compute a sequence of points that get closer and closer to the optimal point. At each step, the algorithm searches a set of points, called a *mesh*, around the *current point*—the point computed at the previous step of the algorithm. The mesh is formed by adding the current point to a scalar multiple of a set of vectors called a *pattern*. If the pattern search algorithm finds a point in the mesh that improves the objective function at the current point, the new point becomes the current point at the next step of the algorithm.

The MADS algorithm is a modification of the GPS algorithm. The algorithms differ in how the set of points forming the mesh is computed. The GPS algorithm uses fixed direction vectors, whereas the MADS algorithm uses a random selection of vectors to define the mesh.

## Performing a Pattern Search

This section provides a brief introduction to using the `patternsearch` functionality from the command line and the Pattern Search Tool, a graphical user interface (GUI) for performing a pattern search. This section covers the following topics:

- “Calling `patternsearch` at the Command Line” on page 3-3
- “Using the Pattern Search Tool” on page 3-3

### Calling `patternsearch` at the Command Line

To perform a pattern search on an unconstrained problem at the command line, you call the function `patternsearch` with the syntax

```
[x fval] = patternsearch(@objfun, x0)
```

where

- `@objfun` is a handle to the objective function.
- `x0` is the starting point for the pattern search.

The results are given by

- `x` — Point at which the final value is attained
- `fval` — Final value of the objective function

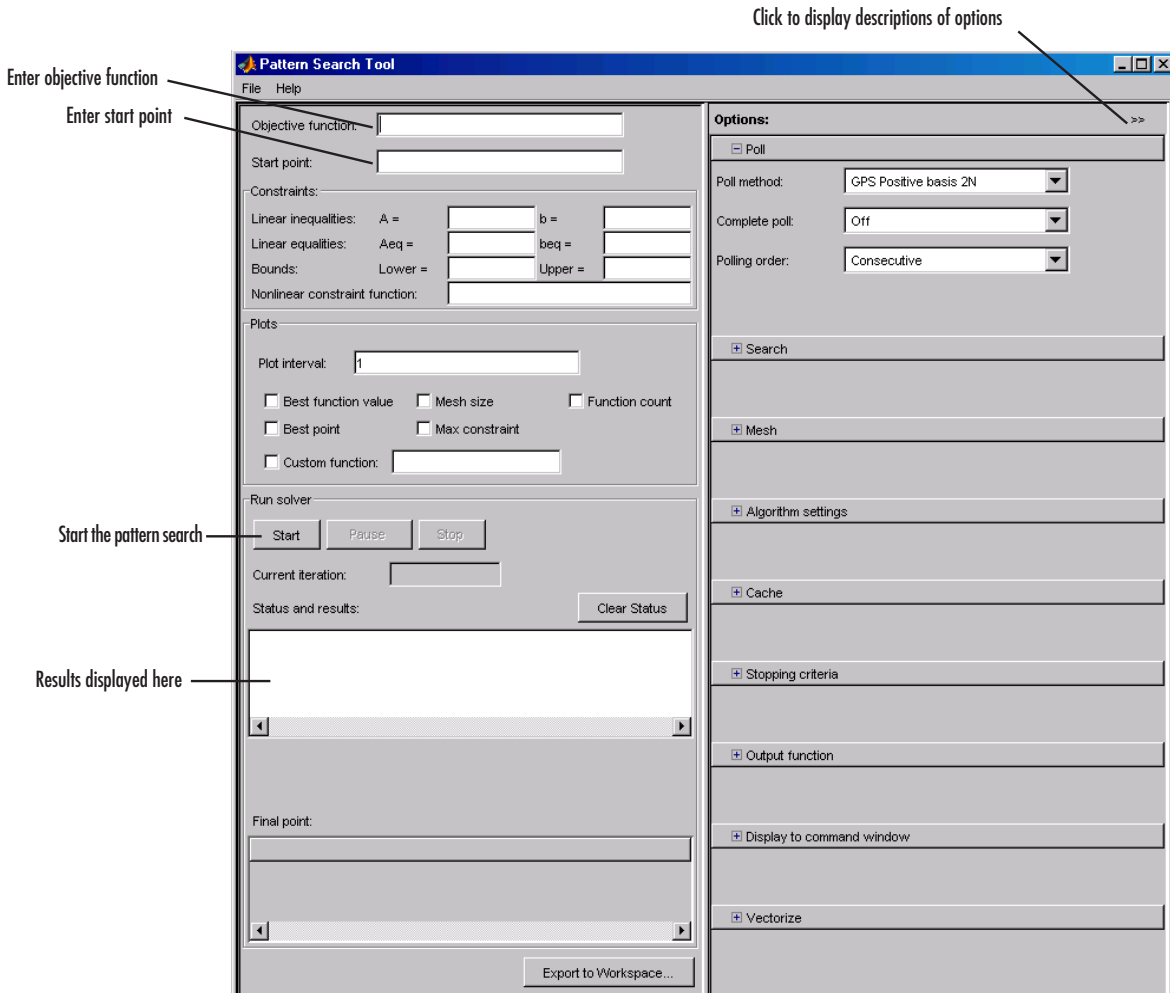
“Performing a Pattern Search from the Command Line” on page 6-20 explains in detail how to use the function `patternsearch`.

### Using the Pattern Search Tool

To open the Pattern Search Tool, enter

```
psearchtool
```

This opens the tool as shown in the following figure.



To use the Pattern Search Tool, you must first enter the following information:

- **Objective function** — The objective function you want to minimize. You enter the objective function in the form `@objfun`, where `objfun.m` is an M-file that computes the objective function. The `@` sign creates a function handle to `objfun`.



- **Start point** — The initial point at which the algorithm starts the optimization.

You can enter linear constraints, bounds, or a nonlinear constraint function for the problem in the **Constraints** pane. If the problem is unconstrained, leave these fields blank.

Then, click the **Start** button. The tool displays the results of the optimization in the **Status and results** pane.

You can also change the options for the pattern search in the **Options** pane. To view the options in a category, click the + sign next to it.

“Finding the Minimum of the Function” on page 3-7 gives an example of using the Pattern Search Tool.

“Overview of the Pattern Search Tool” on page 6-2 provides a detailed description of the Pattern Search Tool.

## Example: Finding the Minimum of a Function Using the GPS Algorithm

This section presents an example of using a pattern search to find the minimum of a function. This section covers the following topics:


- “Objective Function” on page 3-6
- “Finding the Minimum of the Function” on page 3-7
- “Plotting the Objective Function Values and Mesh Sizes” on page 3-8

### Objective Function

The example uses the objective function, `ps_example`, which is included in the Genetic Algorithm and Direct Search Toolbox. You can view the code for the function by entering

```
type ps_example
```

The following figure shows a plot of the function.



## Finding the Minimum of the Function

To find the minimum of `ps_example`, do the following steps:


1 Enter

```
psearchtool
```

to open the Pattern Search Tool.

2 In the **Objective function** field of the Pattern Search Tool, enter `@ps_example`.

3 In the **Start point** field, type `[2.1 1.7]`.




Objective function:

Start point:

You can leave the fields in the **Constraints** pane blank because the problem is unconstrained.

4 Click **Start** to run the pattern search.

The **Status and results** pane displays the results of the pattern search.



Status and results: Clear Status

```
Pattern search running.
Pattern search terminated.
Objective function value: -1.999999237060392
Optimization terminated: mesh size less than option
```

Final point:

| 1        | 2            |
|----------|--------------|
| -4.71239 | -7.62939e-07 |

The reason the optimization terminated is that the mesh size became smaller than the acceptable tolerance value for the mesh size, defined by the **Mesh tolerance** parameter in the **Stopping criteria** panel. The minimum function value is approximately -2. The **Final point** pane displays the point at which the minimum occurs.

## Plotting the Objective Function Values and Mesh Sizes

To see the performance of the pattern search, you can display plots of the best function value and mesh size at each iteration. First, select the following check boxes in the **Plots** pane:

- **Best function value**
- **Mesh size**

Plots

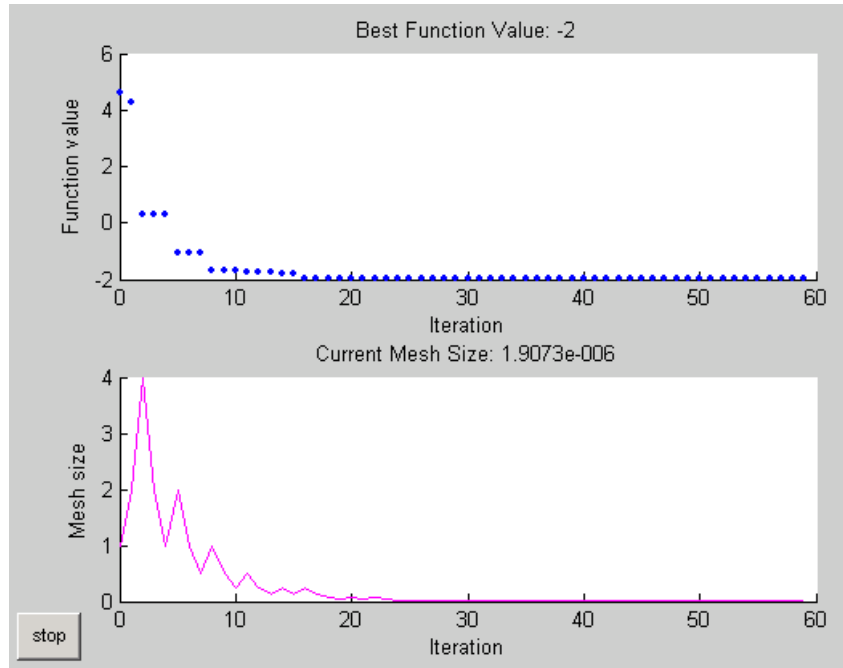
Plot interval:

Best function value     Mesh size     Function count

Best point     Max constraint

Custom function:

Then click **Start** to run the pattern search. This displays the following plots.



The upper plot shows the objective function value of the best point at each iteration. Typically, the objective function values improve rapidly at the early iterations and then level off as they approach the optimal value.

The lower plot shows the mesh size at each iteration. The mesh size increases after each successful iteration and decreases after each unsuccessful one, explained in “How Pattern Search Works” on page 3-14.

## Pattern Search Terminology

This section explains some standard terminology for pattern search, including

- “Patterns” on page 3-10
- “Meshes” on page 3-11
- “Polling” on page 3-12

### Patterns

A *pattern* is a collection of vectors that the pattern search algorithm uses to determine which points to search at each iteration. The set of vectors is defined by the number of independent variables in the objective function,  $N$ , and the positive basis set. Two commonly used positive basis sets in pattern search algorithms are the maximal basis, with  $2N$  vectors, and the minimal basis, with  $N+1$  vectors.


With GPS, the collection of vectors that form the pattern are fixed direction vectors. For example, if there are two independent variables in the optimization problem, the default for a  $2N$  positive basis consists of the following pattern vectors:

$$\begin{aligned}v_1 &= [1 \ 0] \\v_2 &= [0 \ 1] \\v_3 &= [-1 \ 0] \\v_4 &= [0 \ -1]\end{aligned}$$

An  $N+1$  positive basis consists of the following default pattern vectors.

$$\begin{aligned}v_1 &= [1 \ 0] \\v_2 &= [0 \ 1] \\v_3 &= [-1 \ -1]\end{aligned}$$

The following figure shows the pattern vectors for the GPS 2N positive basis and GPS N+1 positive basis.



With MADS, the collection of vectors that form the pattern are randomly selected by the algorithm. Depending on the poll method choice, the number of vectors selected will be 2N or N+1.

## Meshes

At each step, the pattern search algorithm searches a set of points, called a *mesh*, for a point that improves the objective function. Once the direction vectors have been defined, the GPS and MADS algorithms form the mesh by

- 1 Multiplying the pattern vectors by a scalar, called the *mesh size*,  $\Delta^m$ , defined by the length and direction,  $d_i$ , the algorithm moves from the current point,  $x_k$ , to the new point,  $x_i$ .

The new point can be written as the previous point plus the step to the new point  $x_i = x_{i-1} + d_i$

Solving for the mesh size, this is written

$$\Delta^m = \frac{\|x_i - x_{i-1}\|}{\|d_i\|}$$

In GPS, the direction vector is normalized and therefore the mesh size is simply

$$\Delta^m = x_i - x_k$$

- 2** Adding the resulting vectors to the *current point*—the point with the best objective function value found at the previous step.

For example, using the GPS algorithm suppose that

- The current point is [ 1.6 3.4 ].
- The pattern consists of the vectors

$$v_1 = [1 \ 0]$$

$$v_2 = [0 \ 1]$$

$$v_3 = [-1 \ 0]$$

$$v_4 = [0 \ -1]$$

- The current mesh size is 4.

The algorithm multiplies the pattern vectors by 4 and adds them to the current point to obtain the following mesh.

$$[1.6 \ 3.4] + 4*[1 \ 0] = [5.6 \ 3.4]$$

$$[1.6 \ 3.4] + 4*[0 \ 1] = [1.6 \ 7.4]$$

$$[1.6 \ 3.4] + 4*[-1 \ 0] = [-2.4 \ 3.4]$$

$$[1.6 \ 3.4] + 4*[0 \ -1] = [1.6 \ -0.6]$$

The pattern vector that produces a mesh point is called its *direction*.

## Polling

At each step, the algorithm polls the points in the current mesh by computing their objective function values. When option **Complete poll** has the default setting `Off`, the algorithm stops polling the mesh points as soon as it finds a point whose objective function value is less than that of the current point. If this occurs, the poll is called *successful* and the point it finds becomes the current point at the next iteration. Note that the algorithm only computes



the mesh points and their objective function values up to the point at which it stops the poll. If the algorithm fails to find a point that improves the objective function, the poll is called *unsuccessful* and the current point stays the same at the next iteration.

If you set **Complete poll** to On, the algorithm computes the objective function values at all mesh points. The algorithm then compares the mesh point with the smallest objective function value to the current point. If that mesh point has a smaller value than the current point, the poll is successful.

## How Pattern Search Works

The pattern search algorithms find a sequence of points,  $x_0, x_1, x_2, \dots$ , that approaches the optimal point. The value of the objective function either decreases or remains the same from each point in the sequence to the next. This section explains how pattern search works for the function described in “Example: Finding the Minimum of a Function Using the GPS Algorithm” on page 3-6.

To simplify the explanation, this section describes how the generalized pattern search (GPS) works using a maximal positive basis of  $2N$ , with **Scale** set to `Off` in **Mesh** options.

This section covers the following topics:

- “Successful Polls” on page 3-14
- “An Unsuccessful Poll” on page 3-17
- “Displaying the Results at Each Iteration” on page 3-19
- “More Iterations” on page 3-20
- “Stopping Conditions for the Pattern Search” on page 3-21

### Successful Polls


The pattern search begins at the initial point  $x_0$  that you provide. In this example,  $x_0 = [2.1 \ 1.7]$ .

#### Iteration 1

At the first iteration, the mesh size is 1 and the GPS algorithm adds the pattern vectors to the initial point  $x_0 = [2.1 \ 1.7]$  to compute the following mesh points.

$$\begin{aligned} [1 \ 0] + x_0 &= [3.1 \ 1.7] \\ [0 \ 1] + x_0 &= [2.1 \ 2.7] \\ [-1 \ 0] + x_0 &= [1.1 \ 1.7] \\ [0 \ -1] + x_0 &= [2.1 \ 0.7] \end{aligned}$$

The algorithm computes the objective function at the mesh points in the order shown above. The following figure shows the value of `ps_example` at the initial point and mesh points.



The algorithm polls the mesh points by computing their objective function values until it finds one whose value is smaller than 4.6347, the value at  $x_0$ . In this case, the first such point it finds is  $[1.1 \ 1.7]$ , at which the value of the objective function is 4.5146, so the poll at iteration 1 is *successful*. The algorithm sets the next point in the sequence equal to

$$x_1 = [1.1 \ 1.7]$$

---

**Note** By default, the GPS pattern search algorithm stops the current iteration as soon as it finds a mesh point whose fitness value is smaller than that of the current point. Consequently, the algorithm might not poll all the mesh points. You can make the algorithm poll all the mesh points by setting **Complete poll** to On.


---

#### Iteration 2

After a successful poll, the algorithm multiplies the current mesh size by 2, the default value of **Expansion factor** in the **Mesh** options pane. Because the initial mesh size is 1, at the second iteration the mesh size is 2. The mesh at iteration 2 contains the following points.

$$\begin{aligned}2*[1 \ 0] + x_1 &= [3.1 \ 1.7] \\2*[0 \ 1] + x_1 &= [1.1 \ 3.7] \\2*[-1 \ 0] + x_1 &= [-0.9 \ 1.7] \\2*[0 \ -1] + x_1 &= [1.1 \ -0.3]\end{aligned}$$

The following figure shows the point  $x_1$  and the mesh points, together with the corresponding values of  $ps\_example$ .



The algorithm polls the mesh points until it finds one whose value is smaller than 4.5146, the value at  $x_1$ . The first such point it finds is  $[-0.9 \ 1.7]$ , at which the value of the objective function is 3.25, so the poll at iteration 2 is again successful. The algorithm sets the second point in the sequence equal to

$$x_2 = [-0.9 \ 1.7]$$

Because the poll is successful, the algorithm multiplies the current mesh size by 2 to get a mesh size of 4 at the third iteration.

## An Unsuccessful Poll


By the fourth iteration, the current point is

$$x_3 = [-4.9 \ 1.7]$$

and the mesh size is 8, so the mesh consists of the points

$$\begin{aligned} 8*[1 \ 0] + x3 &= [3.1 \ 1.7] \\ 8*[0 \ 1] + x3 &= [-4.9 \ 9.7] \\ 8*[-1 \ 0] + x3 &= [-12.9 \ 1.7] \\ 8*[0 \ -1] + x3 &= [-4.9 \ -1.3] \end{aligned}$$

The following figure shows the mesh points and their objective function values.



At this iteration, none of the mesh points has a smaller objective function value than the value at x3, so the poll is *unsuccessful*. In this case, the algorithm does not change the current point at the next iteration. That is,

$$x4 = x3;$$

At the next iteration, the algorithm multiplies the current mesh size by 0.5, the default value of **Contraction factor** in the **Mesh** options pane, so that the mesh size at the next iteration is 4. The algorithm then polls with a smaller mesh size.

## Displaying the Results at Each Iteration

You can display the results of the pattern search at each iteration by setting **Level of display** to **Iterative** in **Display to command window** options. This enables you to evaluate the progress of the pattern search and to make changes to options if necessary.

With this setting, the pattern search displays information about each iteration at the command line. The first four lines of the display are

| Iter | f-count | f(x)      | MeshSize | Method          |
|------|---------|-----------|----------|-----------------|
| 0    | 1       | 4.63474   | 1        |                 |
| 1    | 4       | 4.51464   | 2        | Successful Poll |
| 2    | 7       | 3.25      | 4        | Successful Poll |
| 3    | 10      | -0.264905 | 8        | Successful Poll |


The entry `Successful Poll` below `Method` indicates that the current iteration was successful. For example, the poll at iteration 2 is successful. As a result, the objective function value of the point computed at iteration 2, displayed below `f(x)`, is less than the value at iteration 1.

At iteration 3, the entry `Refine Mesh` below `Method` tells you that the poll is unsuccessful. As a result, the function value at iteration 4 remains unchanged from iteration 3.

Note that by default the pattern search doubles the mesh size after each successful poll and halves it after each unsuccessful poll.

## More Iterations

The pattern search performs 88 iterations before stopping. The following plot shows the points in the sequence computed in the first 13 iterations of the pattern search.



The numbers below the points indicate the first iteration at which the algorithm finds the point. The plot only shows iteration numbers corresponding to successful polls, because the best point doesn't change after an unsuccessful poll. For example, the best point at iterations 4 and 5 is the same as at iteration 3.



## Stopping Conditions for the Pattern Search

This section describes the criteria for stopping the pattern search algorithm. These criteria are listed in the **Stopping criteria** section of the Pattern Search Tool, as shown in the following figure.

Stopping criteria

Mesh tolerance:

Max iteration:  100\*numberofvariables  
 Specify:

Max function evaluations:  2000\*numberofvariables  
 Specify:

Time limit:

Bind tolerance:

X tolerance:

Function tolerance:

Nonlinear constraint tolerance:

The algorithm stops when any of the following conditions occurs:

- The mesh size is less than **Mesh tolerance**.
- The number of iterations performed by the algorithm reaches the value of **Max iteration**.
- The total number of objective function evaluations performed by the algorithm reaches the value of **Max function evaluations**.
- The time, in seconds, the algorithm runs until it reaches the value of **Time limit**.
- The distance between the point found in two consecutive iterations and the mesh size are both less than **X tolerance**.
- The change in the objective function in two consecutive iterations and the mesh size are both less than **Function tolerance**.

Neither the **Bind tolerance** nor the **Nonlinear constraint tolerance** are used as stopping criterion. The **Bind tolerance** option is used to identify active constraints for constrained problems, and the **Nonlinear constraint tolerance** is used to determine the feasibility with respect to nonlinear constraints.

The MADS algorithm uses the relationship between the mesh size,  $\Delta^m$ , and an additional parameter, called the poll parameter,  $\Delta^p$ , to determine the stopping criteria. For positive basis  $N+1$ , the poll parameter is  $N\sqrt{\Delta^m}$ , and for positive basis  $2N$ , the poll parameter is  $\sqrt{\Delta^m}$ . The relationship for MADS stopping criteria is  $\Delta^m \leq$  **Mesh tolerance**.

## Description of the Nonlinear Constraint Solver

The pattern search algorithm uses the Augmented Lagrangian Pattern Search (ALPS) to solve nonlinear constraint problems. The optimization problem solved by the Augmented Lagrangian Pattern Search (ALPS) algorithm is

$$\text{Minimize}_x f(x)$$

such that

$$\begin{aligned} C_i(x) &\leq 0, i = 1 \dots m \\ C_i(x) &= 0, i = m + 1 \dots mt \\ Ax &\leq b \\ A_{eq}x &= beq \\ LB &\leq x \leq UB \end{aligned}$$

where  $C(x)$  represents the nonlinear inequality and equality constraints,  $m$  is the number of nonlinear inequality constraints, and  $mt$  is the total number of nonlinear constraints.

The Augmented Lagrangian Pattern Search (ALPS) attempts to solve a nonlinear optimization problem with nonlinear constraints, linear constraints, and bounds. In this approach, bounds and linear constraints are handled separately from nonlinear constraints. A subproblem is formulated by combining the objective function and nonlinear constraint function using the Lagrangian and the penalty parameters. A sequence of such optimization problems are approximately minimized using a pattern search algorithm such that the linear constraints and bounds are satisfied.

A subproblem formulation is defined as

$$\Theta(x, \lambda, s, \rho) = f(x) - \sum_{i=1}^m \lambda_i s_i \log(s_i - c_i(x)) + \sum_{i=m+1}^{mt} \lambda_i c_i(x) + \frac{\rho}{2} \sum_{i=m+1}^{mt} c_i(x)^2$$

where the components  $\lambda_i$  of the vector  $\lambda$  are nonnegative and are known as Lagrange multiplier estimates. The elements  $s_i$  of the vector  $s$  are nonnegative

shifts, and  $\rho$  is the positive penalty parameter. The algorithm begins by using an initial value for the penalty parameter (`InitialPenalty`).

The pattern search algorithm minimizes a sequence of the subproblem, which is an approximation of the original problem. When the subproblem is minimized to a required accuracy and satisfies feasibility conditions, the Lagrangian estimates are updated. Otherwise, the penalty parameter is increased by a penalty factor (`PenaltyFactor`). This results in a new subproblem formulation and minimization problem. These steps are repeated until the stopping criteria are met. For a complete description of the algorithm, see the following references:

- Lewis, Robert Michael and Virginia Torczon, “A Globally Convergent Augmented Lagrangian Pattern Search Algorithm for Optimization with General Constraints and Simple Bounds”, *SIAM Journal on Optimization*, Volume 12, Number 4, pages 1075–1089, 2002.
- Conn, A. R., N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Pattern Search Algorithm for Optimization with General Constraints and Simple Bounds”, *SIAM Journal on Numerical Analysis*, Volume 28, Number 2, pages 545–572, 1991.
- Conn, A. R., N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Pattern Search Algorithm for Optimization with General Constraints and Simple Bounds”, *Mathematics of Computation*, Volume 66, Number 217, pages 261–288, 1997.

# Getting Started with the Genetic Algorithm

---

|  |   |
|--|---|
| What Is the Genetic Algorithm?<br>(p. 4-2)               | Introduces the genetic algorithm.   |
| Using the Genetic Algorithm (p. 4-3)                     | Explains how to use the Genetic Algorithm Tool.                                     |
| Example: Rastrigin's Function<br>(p. 4-7)                | Presents an example of solving an optimization problem using the genetic algorithm. |
| Some Genetic Algorithm Terminology (p. 4-15)             | Explains some basic terminology for the genetic algorithm.                          |
| How the Genetic Algorithm Works<br>(p. 4-18)             | Presents an overview of how the genetic algorithm works.                            |
| Description of the Nonlinear Constraint Solver (p. 4-25) | Explains the Augmented Lagrangian Genetic Algorithm (ALGA).                         |

## What Is the Genetic Algorithm?

The genetic algorithm is a method for solving both constrained and unconstrained optimization problems that is based on natural selection, the process that drives biological evolution. The genetic algorithm repeatedly modifies a population of individual solutions. At each step, the genetic algorithm selects individuals at random from the current population to be parents and uses them to produce the children for the next generation. Over successive generations, the population "evolves" toward an optimal solution. You can apply the genetic algorithm to solve a variety of optimization problems that are not well suited for standard optimization algorithms, including problems in which the objective function is discontinuous, nondifferentiable, stochastic, or highly nonlinear.

The genetic algorithm uses three main types of rules at each step to create the next generation from the current population:

- *Selection rules* select the individuals, called *parents*, that contribute to the population at the next generation.
- *Crossover rules* combine two parents to form children for the next generation.
- *Mutation rules* apply random changes to individual parents to form children.

The genetic algorithm differs from a classical, derivative-based, optimization algorithm in two main ways, as summarized in the following table.

| Classical Algorithm  | Genetic Algorithm  |
|--|--|
| Generates a single point at each iteration. The sequence of points approaches an optimal solution. | Generates a population of points at each iteration. The best point in the population approaches an optimal solution. |
| Selects the next point in the sequence by a deterministic computation.                             | Selects the next population by computation which uses random number generators.                                      |

## Using the Genetic Algorithm

There are two ways you can use the genetic algorithm with the toolbox:

- Calling the genetic algorithm function `ga` at the command line.
- Using the Genetic Algorithm Tool, a graphical interface to the genetic algorithm.

This section provides a brief introduction to these methods.

### Calling the Function `ga` at the Command Line

To use the genetic algorithm at the command line, call the genetic algorithm function `ga` with the syntax

```
[x fval] = ga(@fitnessfun, nvars, options)
```

where

- `@fitnessfun` is a handle to the fitness function.
- `nvars` is the number of independent variables for the fitness function.
- `options` is a structure containing options for the genetic algorithm. If you do not pass in this argument, `ga` uses its default options.

The results are given by

- `x` — Point at which the final value is attained
- `fval` — Final value of the fitness function

Using the function `ga` is convenient if you want to

- Return results directly to the MATLAB workspace
- Run the genetic algorithm multiple times with different options, by calling `ga` from an M-file

“Using the Genetic Algorithm from the Command Line” on page 7-24 provides a detailed description of using the function `ga` and creating the options structure.

### **Using the Genetic Algorithm Tool**

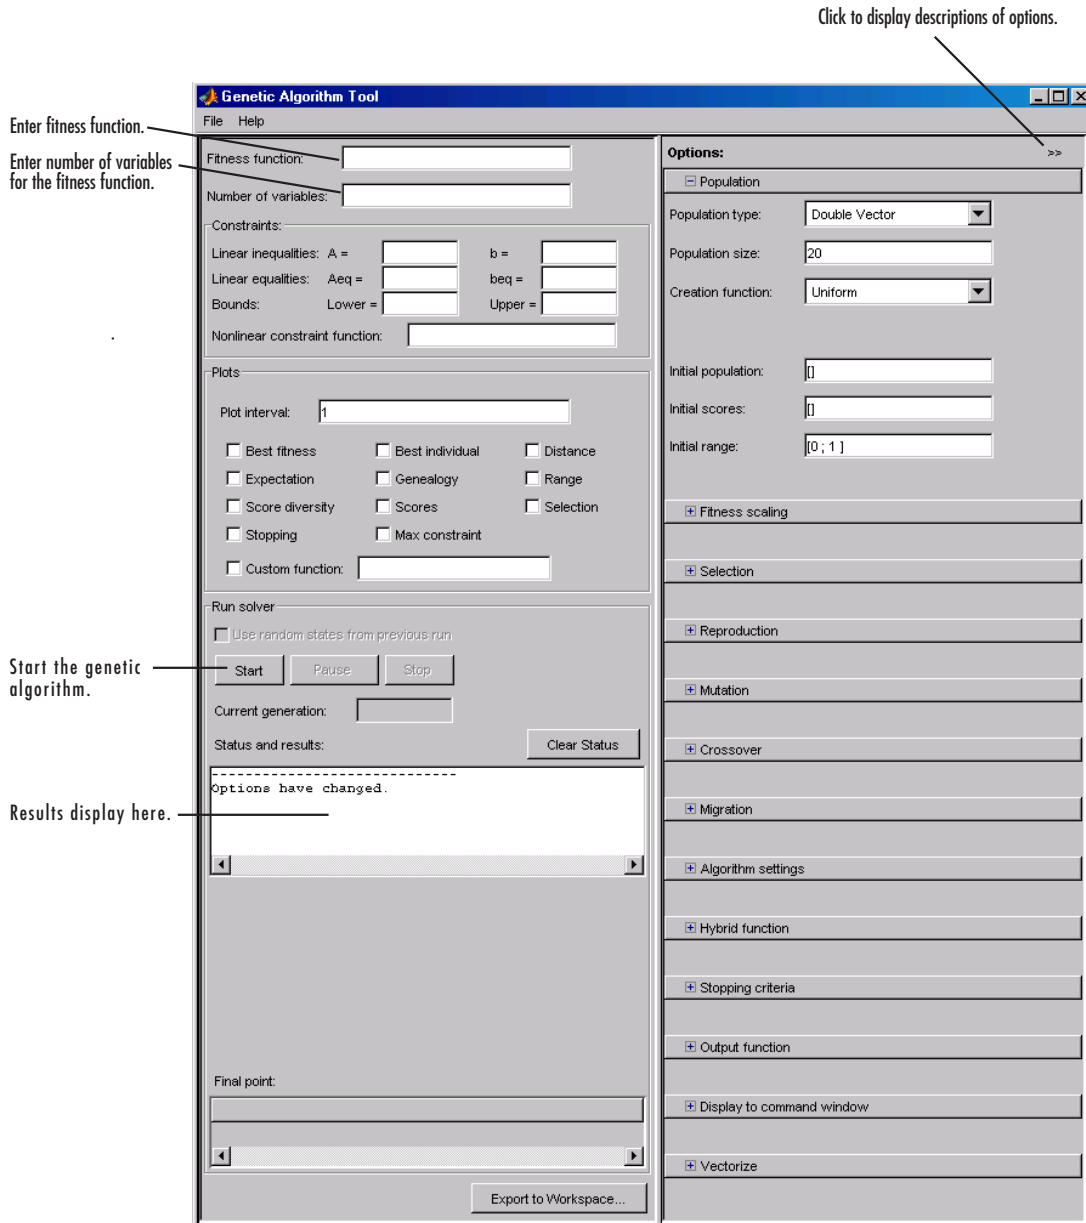
The Genetic Algorithm Tool is a graphical user interface that enables you to use the genetic algorithm without working at the command line. To open the Genetic Algorithm Tool, enter

```
gatool
```

at the MATLAB command prompt.



This opens the tool as shown in the following figure.



To use the Genetic Algorithm Tool, you must first enter the following information:

- **Fitness function** — The objective function you want to minimize. Enter the fitness function in the form `@fitnessfun`, where `fitnessfun.m` is an M-file that computes the fitness function. Chapter 2, “Writing M-Files for Functions You Want to Optimize” explains how to write this M-file. The `@` sign creates a function handle to `fitnessfun`.
- **Number of variables** — The length of the input vector to the fitness function. For the function `my_fun` described in Chapter 2, “Writing M-Files for Functions You Want to Optimize”, you would enter 2.

You can enter constraints or a nonlinear constraint function for the problem in the **Constraints** pane. If the problem is unconstrained, leave these fields blank.

To run the genetic algorithm, click the **Start** button. The tool displays the results of the optimization in the **Status and results** pane.

You can change the options for the genetic algorithm in the **Options** pane. To view the options in one of the categories listed in the pane, click the `+` sign next to it.

For more information,

- See “Overview of the Genetic Algorithm Tool” on page 7-2 for a detailed description of the tool.
- See “Example: Rastrigin’s Function” on page 4-7 for an example of using the tool.

## Example: Rastrigin's Function

This section presents an example that shows how to find the minimum of Rastrigin's function, a function that is often used to test the genetic algorithm. This section covers the following topics:


- “Rastrigin's Function” on page 4-7
- “Finding the Minimum of Rastrigin's Function” on page 4-9
- “Finding the Minimum from the Command Line” on page 4-11
- “Displaying Plots” on page 4-12

### Rastrigin's Function

For two independent variables, Rastrigin's function is defined as

$$Ras(x) = 20 + x_1^2 + x_2^2 - 10(\cos 2\pi x_1 + \cos 2\pi x_2)$$


The toolbox contains an M-file, `rastriginsfcn.m`, that computes the values of Rastrigin's function. The following figure shows a plot of Rastrigin's function.



As the plot shows, Rastrigin's function has many local minima—the “valleys” in the plot. However, the function has just one global minimum, which occurs at the point  $[0\ 0]$  in the  $x$ - $y$  plane, as indicated by the vertical line in the plot, where the value of the function is 0. At any local minimum other than  $[0\ 0]$ , the value of Rastrigin's function is greater than 0. The farther the local minimum is from the origin, the larger the value of the function is at that point.

Rastrigin's function is often used to test the genetic algorithm, because its many local minima make it difficult for standard, gradient-based methods to find the global minimum.

The following contour plot of Rastrigin's function shows the alternating maxima and minima.



## Finding the Minimum of Rastrigin's Function

This section explains how to find the minimum of Rastrigin's function using the genetic algorithm.

---


**Note** Because the genetic algorithm uses random number generators, the algorithm returns slightly different results each time you run it.

---

To find the minimum, do the following steps:


- 1 Enter `gatool` at the command line to open the Genetic Algorithm Tool.
- 2 Enter the following in the Genetic Algorithm Tool:
  - In the **Fitness function** field, enter `@rastriginsfcn`.
  - In the **Number of variables** field, enter 2, the number of independent variables for Rastrigin's function.

The **Fitness function** and **Number of variables** fields should appear as shown in the following figure.



A screenshot of a configuration window showing two input fields. The first field is labeled "Fitness function:" and contains the text "@rastriginsfcn". The second field is labeled "Number of variables:" and contains the number "2".

- 3 Click the **Start** button in the **Run solver** pane, as shown in the following figure.



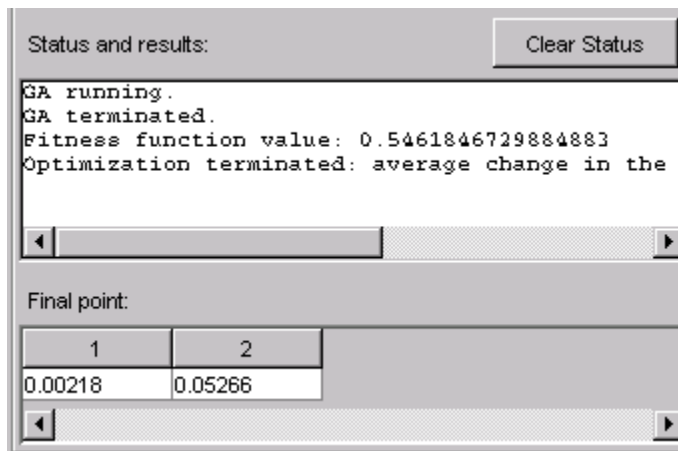
A screenshot of the "Run solver" pane. It features a checkbox labeled "Use random states from previous run" which is unchecked. Below the checkbox are three buttons: "Start", "Pause", and "Stop". A mouse cursor is pointing at the "Start" button. At the bottom of the pane, there is a label "Current generation:" followed by an empty text input field.

While the algorithm is running, the **Current generation** field displays the number of the current generation. You can temporarily pause the

algorithm by clicking the **Pause** button. When you do so, the button name changes to **Resume**. To resume the algorithm from the point at which you paused it, click **Resume**.

When the algorithm is finished, the **Status and results** pane appears as shown in the following figure.

The **Status and results** pane displays the following information:



- The final value of the fitness function when the algorithm terminated:

Function value: 0.5461846729884883

Note that the value shown is very close to the actual minimum value of Rastrigin's function, which is 0. "Genetic Algorithm Examples" on page 7-34 describes some ways to get a result that is closer to the actual minimum.

- The reason the algorithm terminated.

Optimization terminated:  
average change in the fitness value less than options.TolFun.

- The final point, which in this example is [0.00218 0.05266].

## Finding the Minimum from the Command Line

To find the minimum of Rastrigin's function from the command line, enter

```
[x fval exitflag] = ga(@rastriginsfcn, 2)
```

This returns

```
Optimization terminated:  
average change in the fitness value less than options.TolFun.
```

```
x =
```

```
    0.0229    0.0106
```

```
fval =
```

```
    0.1258
```

```
exitflag =
```

```
    1
```

where

- `x` is the final point returned by the algorithm.
- `fval` is the fitness function value at the final point.
- `exitflag` is integer value corresponding to the reason that the algorithm terminated.

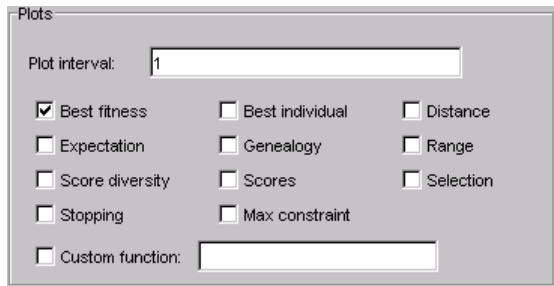
---

**Note** Because the genetic algorithm uses random number generators, the algorithm returns slightly different results each time you run it.


---

### Displaying Plots

The **Plots** pane enables you to display various plots that provide information about the genetic algorithm while it is running. This information can help you change options to improve the performance of the algorithm. For example, to plot the best and mean values of the fitness function at each generation, select the box next to **Best fitness**, as shown in the following figure.



When you click **Start**, the Genetic Algorithm Tool displays a plot of the best and mean values of the fitness function at each generation. When the algorithm stops, the plot appears as shown in the following figure.

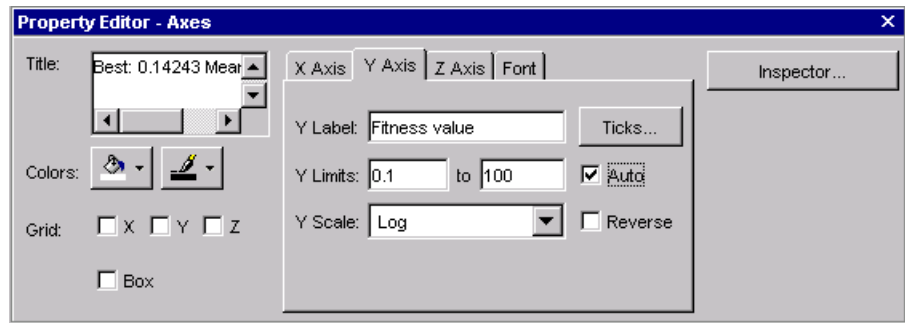




The points at the bottom of the plot denote the best fitness values, while the points above them denote the averages of the fitness values in each generation. The plot also displays the best and mean values in the current generation numerically at the top.


To get a better picture of how much the best fitness values are decreasing, you can change the scaling of the  $y$ -axis in the plot to logarithmic scaling. To do so,

- 1 Select **Axes Properties** from the **Edit** menu in the plot window to open the Property Editor attached to your figure window as shown below.



- 2 Click the **Y** tab.
- 3 In the **Scale** pane, select **Log**.

The plot now appears as shown in the following figure.



Typically, the best fitness value improves rapidly in the early generations, when the individuals are farther from the optimum. The best fitness value improves more slowly in later generations, whose populations are closer to the optimal point.

## Some Genetic Algorithm Terminology

This section explains some basic terminology for the genetic algorithm, including

- “Fitness Functions” on page 4-15
- “Individuals” on page 4-15
- “Populations and Generations” on page 4-16
- “Diversity” on page 4-16
- “Fitness Values and Best Fitness Values” on page 4-16
- “Parents and Children” on page 4-17

### Fitness Functions

The *fitness function* is the function you want to optimize. For standard optimization algorithms, this is known as the objective function. The toolbox tries to find the minimum of the fitness function.

You can write the fitness function as an M-file and pass it as a function handle input argument to the main genetic algorithm function.

### Individuals

An *individual* is any point to which you can apply the fitness function. The value of the fitness function for an individual is its score. For example, if the fitness function is

$$f(x_1, x_2, x_3) = (2x_1 + 1)^2 + (3x_2 + 4)^2 + (x_3 - 2)^2$$

the vector (2, 3, 1), whose length is the number of variables in the problem, is an individual. The score of the individual (2, 3, 1) is  $f(2, -3, 1) = 51$ .

An individual is sometimes referred to as a *genome* and the vector entries of an individual as *genes*.


### Populations and Generations

A *population* is an array of individuals. For example, if the size of the population is 100 and the number of variables in the fitness function is 3, you represent the population by a 100-by-3 matrix. The same individual can appear more than once in the population. For example, the individual (2, 3, 1) can appear in more than one row of the array.

At each iteration, the genetic algorithm performs a series of computations on the current population to produce a new population. Each successive population is called a new *generation*.

### Diversity

*Diversity* refers to the average distance between individuals in a population. A population has high diversity if the average distance is large; otherwise it has low diversity. In the following figure, the population on the left has high diversity, while the population on the right has low diversity.



Diversity is essential to the genetic algorithm because it enables the algorithm to search a larger region of the space.

### Fitness Values and Best Fitness Values

The *fitness value* of an individual is the value of the fitness function for that individual. Because the toolbox finds the minimum of the fitness function,

the *best* fitness value for a population is the smallest fitness value for any individual in the population.

## **Parents and Children**

To create the next generation, the genetic algorithm selects certain individuals in the current population, called *parents*, and uses them to create individuals in the next generation, called *children*. Typically, the algorithm is more likely to select parents that have better fitness values.

# How the Genetic Algorithm Works

This section provides an overview of how the genetic algorithm works. This section covers the following topics:

- “Outline of the Algorithm” on page 4-18
- “Initial Population” on page 4-19
- “Creating the Next Generation” on page 4-20
- “Plots of Later Generations” on page 4-22
- “Stopping Conditions for the Algorithm” on page 4-22

## Outline of the Algorithm


The following outline summarizes how the genetic algorithm works:

- 1 The algorithm begins by creating a random initial population.
- 2 The algorithm then creates a sequence of new populations. At each step, the algorithm uses the individuals in the current generation to create the next population. To create the new population, the algorithm performs the following steps:
  - a Scores each member of the current population by computing its fitness value.
  - b Scales the raw fitness scores to convert them into a more usable range of values.
  - c Selects members, called parents, based on their fitness.
  - d Some of the individuals in the current population that have lower fitness are chosen as *elite*. These elite individuals are passed to the next population.
  - e Produces children from the parents. Children are produced either by making random changes to a single parent—*mutation*—or by combining the vector entries of a pair of parents—*crossover*.
  - f Replaces the current population with the children to form the next generation.

- 3** The algorithm stops when one of the stopping criteria is met. See “Stopping Conditions for the Algorithm” on page 4-22.

## Initial Population

The algorithm begins by creating a random initial population, as shown in the following figure.



In this example, the initial population contains 20 individuals, which is the default value of **Population size** in the **Population** options. Note that all the individuals in the initial population lie in the upper-right quadrant of the picture, that is, their coordinates lie between 0 and 1, because the default value of **Initial range** in the **Population** options is  $[0; 1]$ .

If you know approximately where the minimal point for a function lies, you should set **Initial range** so that the point lies near the middle of that range. For example, if you believe that the minimal point for Rastrigin’s function is near the point  $[0 \ 0]$ , you could set **Initial range** to be  $[-1; 1]$ . However, as this example shows, the genetic algorithm can find the minimum even with a less than optimal choice for **Initial range**.


### Creating the Next Generation

At each step, the genetic algorithm uses the current population to create the children that make up the next generation. The algorithm selects a group of individuals in the current population, called *parents*, who contribute their *genes*—the entries of their vectors—to their children. The algorithm usually selects individuals that have better fitness values as parents. You can specify the function that the algorithm uses to select the parents in the **Selection function** field in the **Selection** options.

The genetic algorithm creates three types of children for the next generation:

- *Elite children* are the individuals in the current generation with the best fitness values. These individuals automatically survive to the next generation.
- *Crossover children* are created by combining the vectors of a pair of parents.
- *Mutation children* are created by introducing random changes, or mutations, to a single parent.

The following schematic diagram illustrates the three types of children.





“Mutation and Crossover” on page 7-43 explains how to specify the number of children of each type that the algorithm generates and the functions it uses to perform crossover and mutation.

The following sections explain how the algorithm creates crossover and mutation children.


### Crossover Children

The algorithm creates crossover children by combining pairs of parents in the current population. At each coordinate of the child vector, the default crossover function randomly selects an entry, or *gene*, at the same coordinate from one of the two parents and assigns it to the child.

### Mutation Children

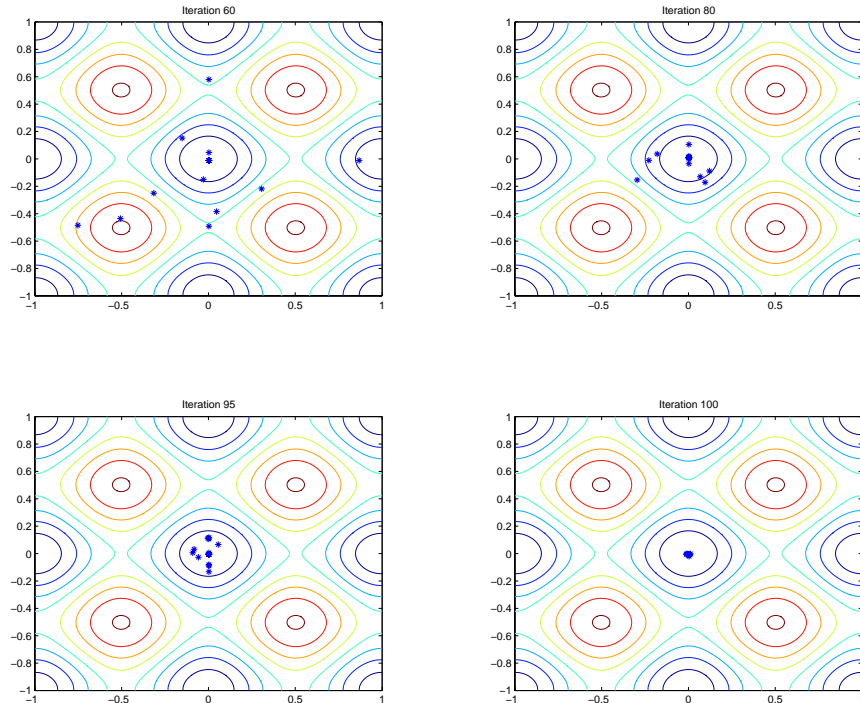
The algorithm creates mutation children by randomly changing the genes of individual parents. By default, the algorithm adds a random vector from a Gaussian distribution to the parent.

The following figure shows the children of the initial population, that is, the population at the second generation, and indicates whether they are elite, crossover, or mutation children.



### Plots of Later Generations

The following figure shows the populations at iterations 60, 80, 95, and 100.



As the number of generations increases, the individuals in the population get closer together and approach the minimum point  $[0 \ 0]$ .

### Stopping Conditions for the Algorithm

The genetic algorithm uses the following conditions to determine when to stop:

- **Generations** — The algorithm stops when the number of generations reaches the value of **Generations**.

- **Time limit** — The algorithm stops after running for an amount of time in seconds equal to **Time limit**.
- **Fitness limit** — The algorithm stops when the value of the fitness function for the best point in the current population is less than or equal to **Fitness limit**.
- **Stall generations** — The algorithm stops when the weighted average change in the fitness function value over **Stall generations** is less than **Function tolerance**.
- **Stall time limit** — The algorithm stops if there is no improvement in the objective function during an interval of time in seconds equal to **Stall time limit**.
- **Function Tolerance** — The algorithm runs until the weighted average change in the fitness function value over **Stall generations** is less than **Function tolerance**.
- **Nonlinear constraint tolerance** — The **Nonlinear constraint tolerance** is not used as stopping criterion. It is used to determine the feasibility with respect to nonlinear constraints.

The algorithm stops as soon as any one of these conditions is met. You can specify the values of these criteria in the **Stopping criteria** pane in the Genetic Algorithm Tool. The default values are shown in the figure below.

| Parameter                       | Value  |
|---------------------------------|--------|
| Generations:                    | 100    |
| Time limit:                     | Inf    |
| Fitness limit:                  | -Inf   |
| Stall generations:              | 50     |
| Stall time limit:               | 20     |
| Function tolerance:             | 1e-006 |
| Nonlinear constraint tolerance: | 1e-006 |

When you run the genetic algorithm, the **Status and results** panel displays the criterion that caused the algorithm to stop.

The options **Stall time limit** and **Time limit** prevent the algorithm from running too long. If the algorithm stops due to one of these conditions, you might improve your results by increasing the values of **Stall time limit** and **Time limit**.

## Description of the Nonlinear Constraint Solver

The genetic algorithm uses the Augmented Lagrangian Genetic Algorithm (ALGA) to solve nonlinear constraint problems. The optimization problem solved by the ALGA algorithm is

$$\text{Minimize}_x f(x)$$

such that

$$\begin{aligned} C_i(x) &\leq 0, i = 1 \dots m \\ C_i(x) &= 0, i = m + 1 \dots mt \\ Ax &\leq b \\ A_{eq}x &= beq \\ LB &\leq x \leq UB \end{aligned}$$

where  $C(x)$  represents the nonlinear inequality and equality constraints,  $m$  is the number of nonlinear inequality constraints, and  $mt$  is the total number of nonlinear constraints.

The Augmented Lagrangian Genetic Algorithm (ALGA) attempts to solve a nonlinear optimization problem with nonlinear constraints, linear constraints, and bounds. In this approach, bounds and linear constraints are handled separately from nonlinear constraints. A subproblem is formulated by combining the fitness function and nonlinear constraint function using the Lagrangian and the penalty parameters. A sequence of such optimization problems are approximately minimized using the genetic algorithm such that the linear constraints and bounds are satisfied.

A subproblem formulation is defined as

$$\Theta(x, \lambda, s, \rho) = f(x) - \sum_{i=1}^m \lambda_i s_i \log(s_i - c_i(x)) + \sum_{i=m+1}^{mt} \lambda_i c_i(x) + \frac{\rho}{2} \sum_{i=m+1}^{mt} c_i(x)^2$$

where the components  $\lambda_i$  of the vector  $\lambda$  are nonnegative and are known as Lagrange multiplier estimates. The elements  $s_i$  of the vector  $s$  are nonnegative

shifts, and  $\rho$  is the positive penalty parameter. The algorithm begins by using an initial value for the penalty parameter (`InitialPenalty`).

The genetic algorithm minimizes a sequence of the subproblem, which is an approximation of the original problem. When the subproblem is minimized to a required accuracy and satisfies feasibility conditions, the Lagrangian estimates are updated. Otherwise, the penalty parameter is increased by a penalty factor (`PenaltyFactor`). This results in a new subproblem formulation and minimization problem. These steps are repeated until the stopping criteria are met. For a complete description of the algorithm, see the following references:

- Conn, A. R., N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Algorithm for Optimization with General Constraints and Simple Bounds,” *SIAM Journal on Numerical Analysis*, Volume 28, Number 2, pages 545–572, 1991.
- Conn, A. R., N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Barrier Algorithm for Optimization with General Inequality Constraints and Simple Bounds,” *Mathematics of Computation*, Volume 66, Number 217, pages 261–288, 1997.

# Getting Started with Simulated Annealing and Threshold Acceptance

---

What Are Simulated Annealing and Threshold Acceptance? (p. 5-2)

Introduces simulated annealing and threshold acceptance.

Calling the Simulated Annealing and Threshold Acceptance Solvers (p. 5-3)

Explains how to use simulated annealing and threshold acceptance.

Example: Minimizing De Jong's Fifth Function (p. 5-4)

Presents an example of solving an optimization problem using simulated annealing.

Some Simulated Annealing and Threshold Acceptance Terminology (p. 5-6)

Explains some basic terminology for simulated annealing and threshold acceptance.

How Simulated Annealing and Threshold Acceptance Work (p. 5-8)

Presents an overview of how the simulated annealing algorithm works.

# What Are Simulated Annealing and Threshold Acceptance?

Simulated annealing is a method for solving unconstrained and bound-constrained optimization problems. The method models the physical process of heating a material and then slowly lowering the temperature to decrease defects, thus minimizing the system energy.

At each iteration of the simulated annealing algorithm, a new point is randomly generated. The distance of the new point from the current point, or the extent of the search, is based on a probability distribution with a scale proportional to the temperature. The algorithm accepts all new points that lower the objective, but also, with a certain probability, points that raise the objective. By accepting points that raise the objective, the algorithm avoids being trapped in local minima, and is able to explore globally for more possible solutions. An *annealing schedule* is selected to systematically decrease the temperature as the algorithm proceeds. As the temperature decreases, the algorithm reduces the extent of its search to converge to a minimum.

Threshold acceptance uses a similar approach, but instead of accepting new points that raise the objective with a certain probability, it accepts all new points below a fixed threshold. The threshold is then systematically lowered, just as the temperature is lowered in an annealing schedule. Because threshold acceptance avoids the probabilistic acceptance calculations of simulated annealing, it may locate an optimizer faster than simulated annealing.



## Calling the Simulated Annealing and Threshold Acceptance Solvers

This section provides a brief introduction to using the simulated annealing and threshold acceptance at the command line. To simplify the presentation, this section uses only the simulated annealing function `simulannealbnd`. The threshold acceptance function `threshacceptbnd` has the same syntax.

To call the simulated annealing function at the command line, use the syntax

```
[x fval] = simulannealbnd(@objfun, x0, lb, ub, options)
```

where

- `@objfun` is a function handle to the objective function.
- `x0` is an initial guess for the optimizer.
- `lb` and `ub` are lower and upper bound constraints, respectively, on `x`.
- `options` is a structure containing options for the algorithm. If you do not pass in this argument, `simulannealbnd` uses its default options.

The results are given by:

- `x` — Final point returned by the solver
- `fval` — Value of the objective function at `x`

The command-line function `simulannealbnd` is convenient if you want to

- Return results directly to the MATLAB workspace.
- Run the simulated annealing algorithm multiple times with different options by calling `simulannealbnd` from an M-file.


“Using the Simulated Annealing and Threshold Acceptance Algorithms from the Command Line” on page 8-2 provides a detailed description of using the function `simulannealbnd` and creating the options structure.

## Example: Minimizing De Jong's Fifth Function

This section presents an example that shows how to find the minimum of the function using simulated annealing.

De Jong's fifth function is a two-dimensional function with many (25) local minima:

```
dejong5fcn
```



Many standard optimization algorithms get stuck in local minima. Because the simulated annealing algorithm performs a wide random search, the chance of being trapped in local minima is decreased.

To run the simulated annealing algorithm without constraints, call `simulannealbnd` at the command line using the objective function in `dejong5fcn.m`:

```
fun = @dejong5fcn;  
[x fval] = simulannealbnd(@fun, [0 0])
```

This returns

```
x =  
-31.9564 -15.9755
```

```
fval =  
5.9288
```

where

- `x` is the final point returned by the algorithm.
- `fval` is the objective function value at the final point.

---

**Note** Because simulated annealing and threshold acceptance both use random number generators, each time you run these algorithms you may get different results. See “Reproducing Your Results” on page 8-5 for more information.

---

# Some Simulated Annealing and Threshold Acceptance Terminology

This section explains some basic terminology for simulated annealing and threshold acceptance, including

- “Objective Function” on page 5-6
- “Temperature” on page 5-6
- “Annealing Schedule” on page 5-6
- “Reannealing” on page 5-6

## Objective Function

The *objective function* is the function you want to optimize. The algorithms in the toolbox attempt to find the minimum of the objective function. Write the objective function as an M-file and pass it to the solver as a function handle.

## Temperature

The *temperature* is the control parameter in simulated annealing that is decreased gradually as the algorithm proceeds. It determines the probability of accepting a worse solution at any step and is used to limit the extent of the search in a given dimension. You can specify the initial temperature as an integer in the `InitialTemperature` option, and the annealing schedule as a function to the `TemperatureFcn` option.

## Annealing Schedule

The *annealing schedule* is the rate by which the temperature is decreased as the algorithm proceeds. The slower the rate of decrease, the better the chances are of finding an optimal solution, but the longer the run time. You can specify the temperature schedule as a function handle with the `TemperatureFcn` option.

## Reannealing

*Annealing* is the technique of closely controlling the temperature when cooling a material to ensure that it is brought to an optimal state. *Reannealing* raises

the temperature after a certain number of new points have been accepted, and starts the search again at the higher temperature. Reannealing avoids getting caught at local minima. You specify the reannealing schedule with the `ReannealInterval` option.

## How Simulated Annealing and Threshold Acceptance Work

- “Outline of the Algorithm” on page 5-8
- “Stopping Conditions for the Algorithm” on page 5-8

### Outline of the Algorithm

The following is an outline of the steps performed for both the simulated annealing and threshold acceptance algorithms:

- 1** The algorithm begins by randomly generating a new point. The distance of the new point from the current point, or the extent of the search, is determined by a probability distribution with a scale proportional to the current temperature.
- 2** The algorithm determines whether the new point is better or worse than the current point. If the new point is better than the current point, it becomes the next point. If the new point is worse than the current point, the algorithm may still make it the next point. Simulated annealing accepts a worse point based on an acceptance probability. Threshold acceptance accepts a worse point if the objective function is raised by less than a fixed threshold.
- 3** The algorithm systematically lowers the temperature and (for threshold acceptance) the threshold, storing the best point found so far.
- 4** Reannealing is performed after a certain number of points (`ReannealInterval`) are accepted by the solver. Reannealing raises the temperature in each dimension, depending on sensitivity information. The search is resumed with the new temperature values.
- 5** The algorithm stops when the average change in the objective function is very small, or when any other stopping criteria are met. See “Stopping Conditions for the Algorithm” on page 5-8.

### Stopping Conditions for the Algorithm

The simulated annealing and threshold acceptance algorithms use the following conditions to determine when to stop:

- `TolFun` — The algorithm runs until the average change in value of the objective function in `StallIterLim` iterations is less than `TolFun`. The default value is  $1e-6$ .
- `MaxIter` — The algorithm stops if the number of iterations exceeds this maximum number of iterations. You can specify the maximum number of iterations as a positive integer or `Inf`. `Inf` is the default.
- `MaxFunEval` specifies the maximum number of evaluations of the objective function. The algorithm stops if the number of function evaluations exceeds the maximum number of function evaluations. The allowed maximum is  $3000 * \text{numberofvariables}$ .
- `TimeLimit` specifies the maximum time in seconds the algorithm runs before stopping.
- `ObjectiveLimit` — The algorithm stops if the best objective function value is less than or equal to the value of `ObjectiveLimit`.





# Using Direct Search

---

|   |   |
|---|---|
| Overview of the Pattern Search Tool (p. 6-2)                | Provides an overview of the Pattern Search Tool.                |
| Performing a Pattern Search from the Command Line (p. 6-20) | Explains how to perform a pattern search from the command line. |
| Pattern Search Examples (p. 6-26)                           | Explains how to set options for a pattern search.               |
| Parameterizing Functions Called by patternsearch (p. 6-53)  | Explains how to write functions with additional parameters.     |

## Overview of the Pattern Search Tool

The section provides an overview of the Pattern Search Tool, the graphical user interface (GUI) for performing a pattern search. This section covers the following topics:

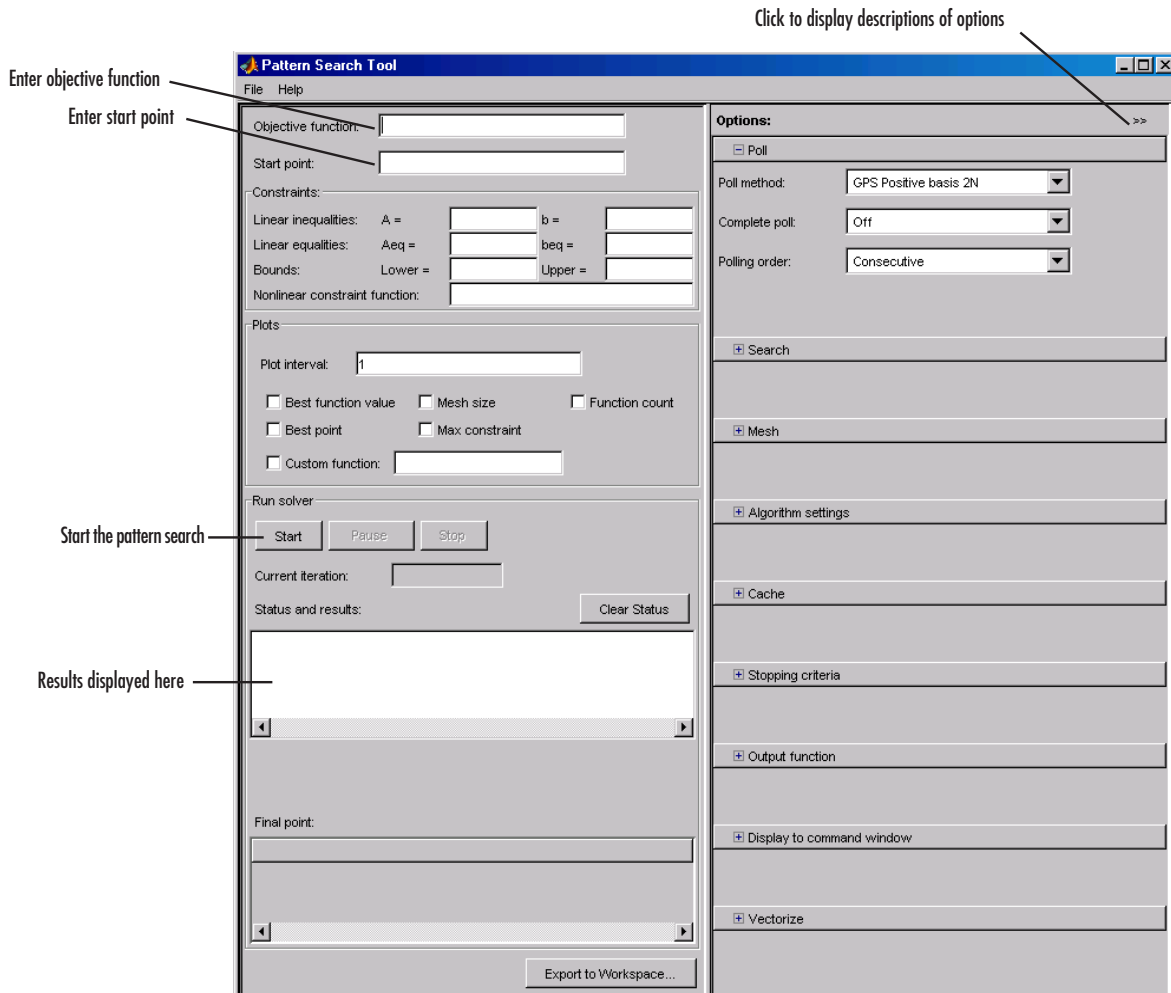
- “Opening the Pattern Search Tool” on page 6-3
- “Defining a Problem in the Pattern Search Tool” on page 6-4
- “Running a Pattern Search” on page 6-5
- “Example — A Linearly Constrained Problem” on page 6-6
- “Pausing and Stopping the Algorithm” on page 6-9
- “Displaying Plots” on page 6-9
- “Example — Working with a Custom Plot Function” on page 6-10
- “Setting Options in the Pattern Search Tool” on page 6-15
- “Importing and Exporting Options and Problems” on page 6-16
- “Generating an M-File” on page 6-19

## Opening the Pattern Search Tool

To open the tool, enter

```
psearchtool
```

at the MATLAB prompt. This opens the Pattern Search Tool, as shown in the following figure.



## Defining a Problem in the Pattern Search Tool

You can define the problem you want to solve using the following fields:

- **Objective function** — The function you want to minimize. Enter a handle to an M-file function that computes the objective function. Chapter 2, “Writing M-Files for Functions You Want to Optimize” describes how to write the M-file.
- **Start point** — The starting point for the pattern search algorithm.

The following figure shows these fields for the example described in “Example: Finding the Minimum of a Function Using the GPS Algorithm” on page 3-6.

The screenshot shows a dialog box for defining a problem. It has the following fields:

- Objective function:
- Start point:
- Constraints:
 

|   |         |                      |         |                      |
|---|---------|----------------------|---------|----------------------|
| Linear inequalities:                                | A =     | <input type="text"/> | b =     | <input type="text"/> |
| Linear equalities:                                  | Aeq =   | <input type="text"/> | beq =   | <input type="text"/> |
| Bounds:   | Lower = | <input type="text"/> | Upper = | <input type="text"/> |
| Nonlinear constraint function: <input type="text"/> |         |                      |         |                      |

## Constrained Problems

You can enter any constraints for the problem in the following fields in the **Constraints** pane:

- **Linear inequalities** — Enter the following for inequality constraints of the form  $Ax \leq b$ :
  - Enter the matrix  $A$  in the **A =** field.
  - Enter the vector  $b$  in the **b =** field.
- **Linear equalities** — Enter the following for equality constraints of the form  $Aeq x = beq$ :
  - Enter the matrix  $Aeq$  in the **Aeq =** field.
  - Enter the vector  $beq$  in the **beq =** field.

- **Bounds** — Enter the following information for bounds constraints of the form  $lb \leq x$  and  $x \leq ub$ :
  - Enter the vector  $lb$  for the lower bound in the **Lower =** field.
  - Enter the vector  $ub$  in the **Upper =** field.

Leave the fields that correspond to constraints, such as the **Nonlinear constraint function** field which does not appear in the problem, empty.

## Running a Pattern Search

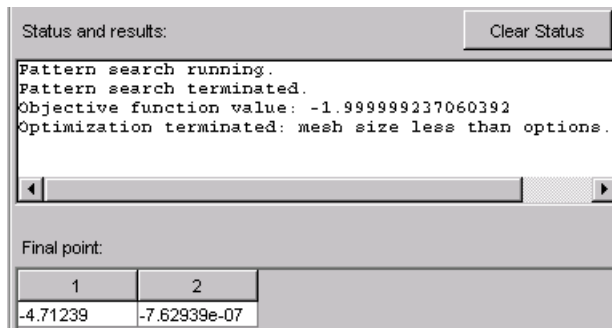
To run a pattern search, click **Start** in the **Run solver** pane. When you do so,

- The **Current iteration** field displays the number of the current iteration.
- The **Status and results** pane displays the message Pattern search running.

When the pattern search terminates, the **Status and results** pane displays

- The message Pattern search terminated.
- The objective function value at the final point
- The reason the pattern search terminated
- The **Final Point** field displays coordinates of the final point.

The following figure shows this information displayed when you run the example in “Example: Finding the Minimum of a Function Using the GPS Algorithm” on page 3-6.



### Example – A Linearly Constrained Problem

This section presents an example of performing a pattern search on a constrained minimization problem. The example minimizes the function

$$F(x) = \frac{1}{2}x^T Hx + f^T x$$

where

$$H = \begin{bmatrix} 36 & 17 & 19 & 12 & 8 & 15 \\ 17 & 33 & 18 & 11 & 7 & 14 \\ 19 & 18 & 43 & 13 & 8 & 16 \\ 12 & 11 & 13 & 18 & 6 & 11 \\ 8 & 7 & 8 & 6 & 9 & 8 \\ 15 & 14 & 16 & 11 & 8 & 29 \end{bmatrix}$$

$$f = [20 \ 15 \ 21 \ 18 \ 29 \ 24]$$

subject to the constraints

$$A \cdot x \leq b$$

$$Aeq \cdot x = beq$$

where

$$A = [-8 \ 7 \ 3 \ -4 \ 9 \ 0]$$

$$b = [7]$$

$$Aeq = \begin{bmatrix} 7 & 1 & 8 & 3 & 3 & 3 \\ 5 & 0 & 5 & 1 & 5 & 8 \\ 2 & 6 & 7 & 1 & 1 & 8 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$beq = [84 \ 62 \ 65 \ 1]$$

## Performing a Pattern Search on the Example

To perform a pattern search on the example, first enter

```
psearchtool
```

to open the Pattern Search Tool. Then type the following function in the **Objective function** field:

```
@lincontest7
```

This is an M-file included in the toolbox that computes the objective function for the example. Because the matrices and vectors defining the starting point and constraints are large, it is more convenient to set their values as variables in the MATLAB workspace first and then enter the variable names in the Pattern Search Tool. To do so, enter

```
x0 = [2 1 0 9 1 0];  
Aineq = [-8 7 3 -4 9 0];  
bineq = [7];  
Aeq = [7 1 8 3 3 3; 5 0 5 1 5 8; 2 6 7 1 1 8; 1 0 0 0 0 0];  
beq = [84 62 65 1];
```

Then, enter the following in the Pattern Search Tool:

- Set **Initial point** to `x0`.
- Set the following **Linear inequalities**:
  - Set **A** = to `Aineq`.
  - Set **b** = to `bineq`.
  - Set **Aeq** = to `Aeq`.
  - Set **beq** = to `beq`.

The following figure shows these settings in the Pattern Search Tool.

Objective function:

Start point:

Constraints:

Linear inequalities: A =  b =

Linear equalities: Aeq =  beq =

Bounds: Lower =  Upper =

Nonlinear constraint function:

Then click **Start** to run the pattern search. When the search is finished, the results are displayed in **Status and results** pane, as shown in the following figure.

Run solver

Current iteration:

Status and results:

```

-----
Pattern search running.
Pattern search terminated.
Objective function value: 2189.0300382498413
Optimization terminated: mesh size less than options.
    
```

Final point:

|  | 1     | 2        | 3       | 4        | 5        |
|--|-------|----------|---------|----------|----------|
|  | 1.001 | -2.30269 | 9.51314 | -0.04738 | -0.19772 |




## Pausing and Stopping the Algorithm

While pattern search is running, you can

- Click **Pause** to temporarily suspend the algorithm. To resume the algorithm from the point at which you paused, click **Resume**.
- Click **Stop** to stop the algorithm. The **Status and results** pane displays the objective function value of the current point at the moment you clicked **Stop**.


## Displaying Plots

The **Plots** pane, shown in the following figure, enables you to display various plots of the results of a pattern search.



The screenshot shows a window titled "Plots" with a light gray background. At the top left, the title "Plots" is displayed. Below the title, there is a text label "Plot interval:" followed by a text input field containing the number "1". Below this, there are five checkboxes arranged in two rows. The first row contains three checkboxes: "Best function value" (checked), "Mesh size" (checked), and "Function count" (unchecked). The second row contains two checkboxes: "Best point" (unchecked) and "Max constraint" (unchecked). Below the second row, there is a checkbox labeled "Custom function:" followed by an empty text input field.

Select the check boxes next to the plots you want to display. For example, if you select **Best function value** and **Mesh size**, and run the example described in “Example: Finding the Minimum of a Function Using the GPS Algorithm” on page 3-6, the tool displays the plots shown in the following figure.



The upper plot displays the objective function value at each iteration. The lower plot displays the mesh size at each iteration.

---

**Note** When you display more than one plot, clicking on any plot while the pattern search is running or after the solver has completed opens a larger version of the plot in a separate window.

---

“Plot Options” on page 9-24 describes the types of plots you can create.

### Example – Working with a Custom Plot Function

To use a plot function other than those included with the toolbox, you can write your own custom plot function that is called at each iteration of the pattern search to create the plot. This example shows how to create a plot function that displays the logarithmic change in the best objective function value from the previous iteration to the current iteration.

This section covers the following topics:

- “Creating the Custom Plot Function” on page 6-11
- “Using the Custom Plot Function” on page 6-12
- “How the Plot Function Works” on page 6-13

## Creating the Custom Plot Function

To create the plot function for this example, copy and paste the following code into a new M-file in the MATLAB Editor.

```
function stop = psplotchange(optimvalues, flag)
% PSLOTCHANGE Plots the change in the best objective function
% value from the previous iteration.

% Best objective function value in the previous iteration
persistent last_best

stop = false;
if(strcmp(flag,'init'))
    set(gca,'Yscale','log'); %Set up the plot
    hold on;
    xlabel('Iteration');
    ylabel('Log Change in Values');
    title(['Change in Best Function Value']);
end

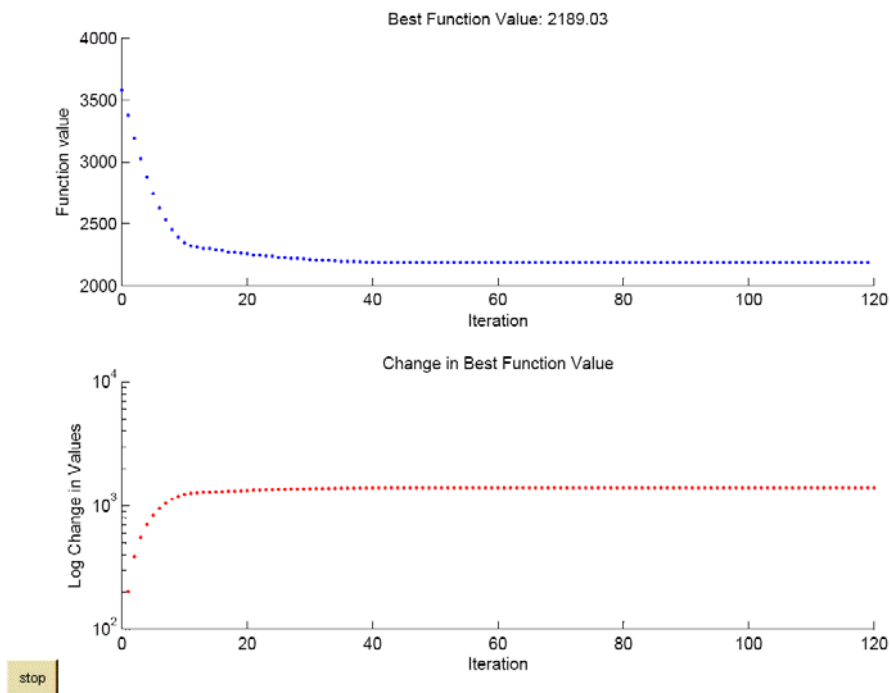
% Best objective function value in the current iteration
best = min(optimvalues.fval);

% Set last_best to best
if optimvalues.iteration == 0
    last_best = best;
else
    %Change in objective function value
    change = last_best - best;
    plot(optimvalues.iteration, change, '.r');
end
```

Then save the M-file as `psplotchange.m` in a directory on the MATLAB path.

### Using the Custom Plot Function

To use the custom plot function, select **Custom function** in the **Plots** pane and enter `@psplotchange` in the field to the right. To compare the custom plot with the best function value plot, also select **Best function value**. Now, when you run the example described in “Example — A Linearly Constrained Problem” on page 6-6, the pattern search tool displays the plots shown in the following figure.



Note that because the scale of the y-axis in the lower custom plot is logarithmic, the plot will only show changes that are greater than 0. The

logarithmic scale shows small changes in the objective function that the upper plot does not reveal.

### How the Plot Function Works

The plot function uses information contained in the following structures, which the Pattern Search Tool passes to the function as input arguments:

- `optimvalues` — Structure containing the current state of the solver
- `flag` — String indicating the current status of the algorithm

The most important statements of the custom plot function, `psplotchange.m`, are summarized in the following table.

### Custom Plot Function Statements

| M-File Statement                          | Description   |
|---|---|
| <code>persistent last_best</code>         | Creates the persistent variable <code>last_best</code> , the best objective function value in the previous generation. Persistent variables are preserved over multiple calls to the plot function.   |
| <code>set(gca, 'Yscale', 'log')</code>    | Sets up the plot before the algorithm starts.   |
| <code>best = min(optimvalues.fval)</code> | Sets <code>best</code> equal to the minimum objective function value. The field <code>optimvalues.fval</code> contains the objective function value in the current iteration. The variable <code>best</code> is the minimum objective function value. For a complete description of the fields of the structure <code>optimvalues</code> , see “Structure of the Plot Functions” on page 9-4. |

**Custom Plot Function Statements (Continued)**

| <b>M-File Statement</b>                                | <b>Description</b>  |
|--|---|
| <code>change = last_best - best</code>                 | Sets the variable change to the best objective function value at the previous iteration minus the best objective function value in the current iteration. |
| <code>plot(optimvalues.iteration, change, '.r')</code> | Plots the variable change at the current objective function value, for the current iteration contained in <code>optimvalues.iteration</code> .            |

## Setting Options in the Pattern Search Tool

You can set options for a pattern search in the **Options** pane, shown in the figure below.

**Options:** >>

Poll

Poll method: GPS Positive basis 2N ▼

Complete poll: Off ▼

Polling order: Consecutive ▼

Search

Mesh

Algorithm settings

Cache

Stopping criteria

Output function

Display to command window

Vectorize

For a detailed description of the available options, see .

### **Setting Options as Variables in the MATLAB Workspace**

You can set numerical options either directly, by typing their values in the corresponding edit box, or by entering the name of a variable in the MATLAB workspace that contains the option values. For options whose values are large matrices or vectors, it is often more convenient to define their values as variables in the MATLAB workspace.

### **Importing and Exporting Options and Problems**

You can export options and problem structures from the Pattern Search Tool to the MATLAB workspace, and later import them in a subsequent session of the tool. This provides an easy way to save your work for future sessions of the Pattern Search Tool. The following sections describe how to import and export options and problem structures.

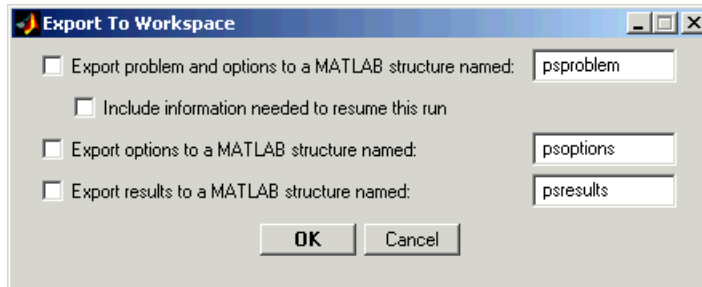
#### **Exporting Options, Problems, and Results**

After working on a problem using the Pattern Search Tool, you can export the following information to the MATLAB workspace:

- The problem definition, including
  - The objective function
  - The start point
  - Constraints on the problem
- The current options
- The results of the algorithm



To do so, click the **Export to Workspace** button or select **Export to Workspace** from the **File** menu. This opens the dialog box shown in the following figure.



The dialog box provides the following options:

- To save the objective function and options in a MATLAB structure, select **Export problem and options to a MATLAB structure named** and enter a name for the structure.

If you have run a pattern search in the current session and you select **Include information needed to resume this run**, the final point from the last search is saved in place of **Start point**. Use this option if you want to run the pattern search at a later time from the final point of the last search.

See “Importing a Problem” on page 6-19.

- To save only the options, select **Export options to a MATLAB structure named** and enter a name for the options structure.
- To save the results of the last run of the algorithm, select **Export results to a MATLAB structure named** and enter a name for the results structure.

### Example — Running patternsearch on an Exported Problem

To export the problem described in “Example — A Linearly Constrained Problem” on page 6-6 and perform a pattern search on it using the function patternsearch at the command line, do the following steps:

- 1 Click **Export to Workspace**.

- 2 In the Export to Workspace dialog box, enter a name for the problem structure, such as `my_psproblem`, in the **Export problems and options to a MATLAB structure named** field.
- 3 Call the function `patternsearch` with `my_psproblem` as the input argument.

```
[x fval] = patternsearch(my_psproblem)
```

This returns

```
x =
```

```
1.0010 -2.3027 9.5131 -0.0474 -0.1977 1.3083
```

```
fval =
```

```
2.1890e+003
```

See “Performing a Pattern Search from the Command Line” on page 6-20 for more information.

### Importing Options

To import an options structure for a pattern search from the MATLAB workspace, select **Import Options** from the **File** menu. This opens a dialog box that displays a list of the valid pattern search options structures in the MATLAB workspace. When you select an options structure and click **Import**, the Pattern Search Tool resets its options to the values in the imported structure.

---

**Note** You cannot import options structures that contain any invalid option fields. Structures with invalid fields are not displayed in the Import Pattern Search Options dialog box.

---

You can create an options structure in either of the following ways:

- Calling `psoptimset` with options as the output
- By saving the current options from the Export to Workspace dialog box in the Pattern Search Tool

### Importing a Problem

To import a problem that you previously exported from the Pattern Search Tool, select **Import Problem** from the **File** menu. This opens the dialog box that displays a list of the pattern search problem structures in the MATLAB workspace. When you select a problem structure and click **OK**, the Pattern Search Tool resets the problem definition and the options to the values in the imported structure. In addition, if you selected **Include information needed to resume this run** when you created the problem structure, the tool resets **Start point** to the final point of the last run prior to exporting the structure.

See “Exporting Options, Problems, and Results” on page 6-16.

### Generating an M-File

To create an M-file that runs a pattern search using the objective function and options you specify in the Pattern Search Tool, select **Generate M-File** from the **File** menu and save the M-file in a directory on the MATLAB path. Calling this M-file at the command line returns the same results as the Pattern Search Tool, using the objective function and options settings that were in place when you generated the M-file.

## Performing a Pattern Search from the Command Line

As an alternative to using the Pattern Search Tool, you can call the function `patternsearch` at the command line. This section explains how to do so and covers the following topics:

- “Calling `patternsearch` with the Default Options” on page 6-20
- “Setting Options for `patternsearch` at the Command Line” on page 6-22
- “Using Options and Problems from the Pattern Search Tool” on page 6-24

### Calling `patternsearch` with the Default Options

This section describes how to perform a pattern search with the default options.

#### Pattern Search on Unconstrained Problems

For an unconstrained problem, call `patternsearch` with the syntax

```
[x fval] = patternsearch(@objectfun, x0)
```

The output arguments are

- `x` — The final point
- `fval` — The value of the objective function at `x`

The required input arguments are

- `@objectfun` — A function handle to the objective function `objectfun`, which you can write as an M-file. See Chapter 2, “Writing M-Files for Functions You Want to Optimize” to learn how to do this.
- `x0` — The initial point for the pattern search algorithm.

As an example, you can run the example described in “Example: Finding the Minimum of a Function Using the GPS Algorithm” on page 3-6 from the command line by entering

```
[x fval] = patternsearch(@ps_example, [2.1 1.7])
```

This returns

```
Optimization terminated: mesh size less than options.TolMesh.
```

```
x =
```

```
    -4.7124    -0.0000
```

```
fval =
```

```
    -2.0000
```

### Pattern Search on Constrained Problems

If your problem has constraints, use the syntax

```
[x fval] = patternsearch(@objfun,x0,A,b,Aeq,beq,lb,ub,nonlcon)
```

where

- $A$  is a matrix and  $b$  is vector that represent inequality constraints of the form  $Ax \leq b$ .
- $Aeq$  is a matrix and  $beq$  is a vector that represent equality constraints of the form  $Aeq x = beq$ .
- $lb$  and  $ub$  are vectors representing bound constraints of the form  $lb \leq x$  and  $x \leq ub$ , respectively.
- $nonlcon$  is a function that returns the nonlinear equality and inequality vectors,  $C$  and  $Ceq$ , respectively. The function is minimized such that  $C(x) \leq 0$  and  $Ceq(x)=0$ .

You only need to pass in the constraints that are part of the problem. For example, if there are no bound constraints or a nonlinear constraint function, use the syntax

```
[x fval] = patternsearch(@objfun, x0, A, b Aeq, beq)
```

Use empty brackets `[]` for constraint arguments that are not needed for the problem. For example, if there are no inequality constraints or a nonlinear constraint function, use the syntax

```
[x fval] = patternsearch(@objfun, x0, [], [], Aeq, beq, lb, ub)
```

### **Additional Output Arguments**

To get more information about the performance of the pattern search, you can call `patternsearch` with the syntax

```
[x fval exitflag output] = patternsearch(@objfun, x0)
```

Besides `x` and `fval`, this returns the following additional output arguments:

- `exitflag` — Integer indicating whether the algorithm was successful
- `output` — Structure containing information about the performance of the solver

See the reference page for `patternsearch` for more information about these arguments.

### **Setting Options for `patternsearch` at the Command Line**

You can specify any of the options that are available in the Pattern Search Tool by passing an options structure as an input argument to `patternsearch` using the syntax

```
[x fval] = patternsearch(@fitnessfun, nvars, ...  
                        A, b, Aeq, beq, lb, ub, nonlcon, options)
```

Pass in empty brackets `[]` for any constraints that do not appear in the problem.

You create the options structure using the function `psoptimset`.

```
options = psoptimset
```

This returns the options structure with the default values for its fields.

```
options =  
  
    TolMesh: 1.0000e-006  
    TolCon: 1.0000e-006  
    TolX: 1.0000e-006  
    TolFun: 1.0000e-006  
    TolBind: 1.0000e-003  
    MaxIter: '100*numberofvariables'  
    MaxFunEvals: '2000*numberofvariables'  
    TimeLimit: Inf  
MeshContraction: 0.5000  
    MeshExpansion: 2  
MeshAccelerator: 'off'  
    MeshRotate: 'on'  
InitialMeshSize: 1  
    ScaleMesh: 'on'  
    MaxMeshSize: Inf  
InitialPenalty: 10  
PenaltyFactor: 100  
    PollMethod: 'gpspositivebasis2n'  
    CompletePoll: 'off'  
    PollingOrder: 'consecutive'  
    SearchMethod: []  
CompleteSearch: 'off'  
    Display: 'final'  
    OutputFcns: []  
    PlotFcns: []  
PlotInterval: 1  
    Cache: 'off'  
    CacheSize: 10000  
    CacheTol: 2.2204e-016  
    Vectorized: 'off'
```

The function `patternsearch` uses these default values if you do not pass in `options` as an input argument.

The value of each option is stored in a field of the `options` structure, such as `options.MeshExpansion`. You can display any of these values by entering `options` followed by the name of the field. For example, to display the mesh expansion factor for the pattern search, enter

```
options.MeshExpansion
```

```
ans =
```

```
2
```

To create an options structure with a field value that is different from the default, use the function `psoptimset`. For example, to change the mesh expansion factor to 3 instead of its default value 2, enter

```
options = psoptimset('MeshExpansion', 3)
```

This creates the options structure with all values set to their defaults except for `MeshExpansion`, which is set to 3.

If you now call `patternsearch` with the argument `options`, the pattern search uses a mesh expansion factor of 3.

If you subsequently decide to change another field in the options structure, such as setting `PlotFcns` to `@psplotmeshsize`, which plots the mesh size at each iteration, call `psoptimset` with the syntax

```
options = psoptimset(options, 'PlotFcns', @psplotmeshsize)
```

This preserves the current values of all fields of options except for `PlotFcns`, which is changed to `@plotmeshsize`. Note that if you omit the options input argument, `psoptimset` resets `MeshExpansion` to its default value, which is 2.

You can also set both `MeshExpansion` and `PlotFcns` with the single command

```
options = psoptimset('MeshExpansion',3,'PlotFcns',@plotmeshsize)
```

## Using Options and Problems from the Pattern Search Tool

As an alternative to creating the options structure using `psoptimset`, you can set the values of options in the Pattern Search Tool and then export the options to a structure in the MATLAB workspace, as described in “Exporting Options, Problems, and Results” on page 6-16. If you export the default options in the Pattern Search Tool, the resulting options structure has the same settings as the default structure returned by the command



```
options = psoptimset
```

except for the default value of 'Display', which is 'final' when created by psoptimset, but 'none' when created in the Pattern Search Tool.

You can also export an entire problem from the Pattern Search Tool and run it from the command line. See “Example — Running patternsearch on an Exported Problem” on page 6-17 for an example.

## Pattern Search Examples

This section explains how to set options for a pattern search. All examples use the generalized pattern search (GPS) algorithm, but can be applied to the MADS algorithm as well.

- “Poll Method” on page 6-26
- “Complete Poll” on page 6-28
- “Using a Search Method” on page 6-32
- “Mesh Expansion and Contraction” on page 6-35
- “Mesh Accelerator” on page 6-40
- “Using Cache” on page 6-41
- “Setting Tolerances for the Solver” on page 6-44
- “Constrained Minimization Using patternsearch” on page 6-50

### Poll Method

At each iteration, the pattern search polls the points in the current mesh—that is, it computes the objective function at the mesh points to see if there is one whose function value is less than the function value at the current point. “How Pattern Search Works” on page 3-14 provides an example of polling. You can specify the pattern that defines the mesh by the **Poll method** option. The default pattern, GPS Positive basis  $2N$ , consists of the following  $2N$  directions, where  $N$  is the number of independent variables for the objective function.

```
[100...0]
[010...0]
...
[000...1]
[-1 00...0]
[0 -10...0]
...
[0 00...-1]
```

For example, if objective function has three independent variables, the GPS Positive basis  $2N$ , consists of the following six vectors.

$$\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & -1 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 & -1 \end{bmatrix}$$

Alternatively, you can set **Poll method** to GPS Positive basis  $N+1$ , the pattern consisting of the following  $N + 1$  directions.

$$\begin{bmatrix} 1 & 0 & \dots & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1 & \dots & 0 \end{bmatrix}$$

$$\dots$$

$$\begin{bmatrix} 0 & 0 & \dots & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & -1 & \dots & -1 \end{bmatrix}$$

For example, if objective function has three independent variables, the GPS Positive basis  $N+1$ , consists of the following four vectors.

$$\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & -1 & -1 \end{bmatrix}$$

A pattern search will sometimes run faster using GPS Positive basis  $N+1$  rather than the GPS Positive basis  $2N$  as the **Poll method**, because the algorithm searches fewer points at each iteration. Although not being addressed in this example, the same is true when using the MADS Positive basis  $N+1$  over the MADS Positive basis  $2N$ . For example, if you run a pattern search on the example described in “Example — A Linearly

Constrained Problem” on page 6-6, the algorithm performs 2080 function evaluations with GPS Positive basis 2N, the default **Poll method**, but only 1413 function evaluations using GPS Positive basis Np1.

However, if the objective function has many local minima, using GPS Positive basis 2N as the **Poll method** might avoid finding a local minimum that is not the global minimum, because the search explores more points around the current point at each iteration.

## Complete Poll

By default, if the pattern search finds a mesh point that improves the value of the objective function, it stops the poll and sets that point as the current point for the next iteration. When this occurs, some mesh points might not get polled. Some of these unpolled points might have an objective function value that is even lower than the first one the pattern search finds.

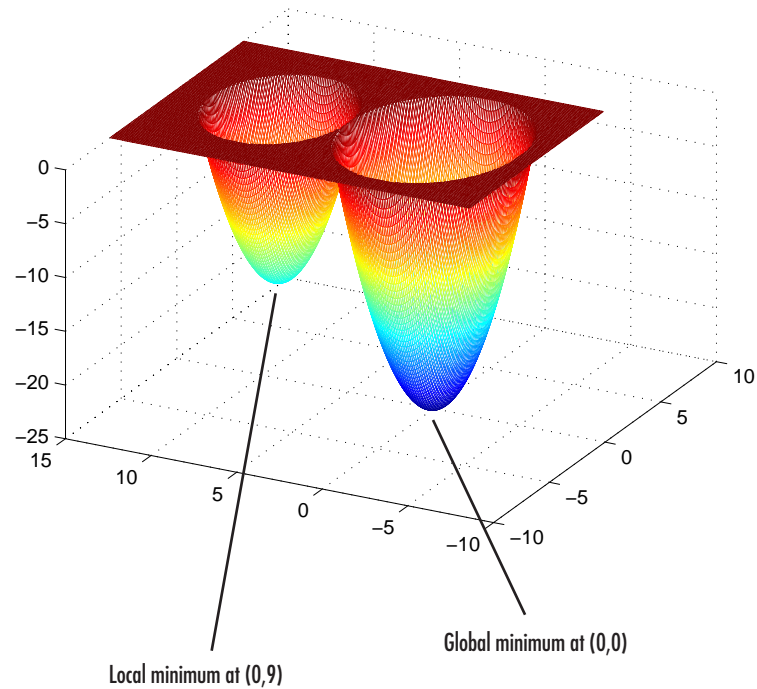
For problems in which there are several local minima, it is sometimes preferable to make the pattern search poll *all* the mesh points at each iteration and choose the one with the best objective function value. This enables the pattern search to explore more points at each iteration and thereby potentially avoid a local minimum that is not the global minimum. You can make the pattern search poll the entire mesh setting **Complete poll** to On in **Poll** options.

## Example – Using a Complete Poll in a Generalized Pattern Search

As an example, consider the following function.

$$f(x_1, x_2) = \begin{cases} x_1^2 + x_2^2 - 25 & \text{for } x_1^2 + x_2^2 \leq 25 \\ x_1^2 + (x_2 - 9)^2 - 16 & \text{for } x_1^2 + (x_2 - 9)^2 \leq 16 \\ 0 & \text{otherwise} \end{cases}$$

The following figure shows a plot of the function.



The global minimum of the function occurs at (0, 0), where its value is -25. However, the function also has a local minimum at (0, 9), where its value is -16.

To create an M-file that computes the function, copy and paste the following code into a new M-file in the MATLAB Editor.

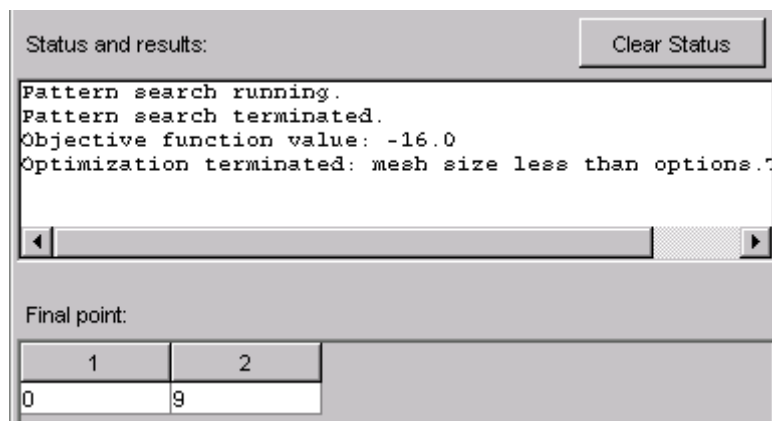
```
function z = poll_example(x)
if x(1)^2 + x(2)^2 <= 25
    z = x(1)^2 + x(2)^2 - 25;
elseif x(1)^2 + (x(2) - 9)^2 <= 16
    z = x(1)^2 + (x(2) - 9)^2 - 16;
else z = 0;
end
```

Then save the file as `poll_example.m` in a directory on the MATLAB path.

To run a pattern search on the function, enter the following in the Pattern Search Tool:

- Set **Objective function** to `@poll_example`.
- Set **Start point** to `[0 5]`.
- Set **Level of display** to **Iterative** in the **Display to command window** options.

Click **Start** to run the pattern search with **Complete poll** set to **Off**, its default value. The Pattern Search Tool displays the results in the **Status and results** pane, as shown in the following figure.



The pattern search returns the local minimum at (0, 9). At the initial point, (0, 5), the objective function value is 0. At the first iteration, the search polls the following mesh points.

$$f(0, 5) + (1, 0) = f(1, 5) = 0$$

$$f(0, 5) + (0, 1) = f(0, 6) = -7$$

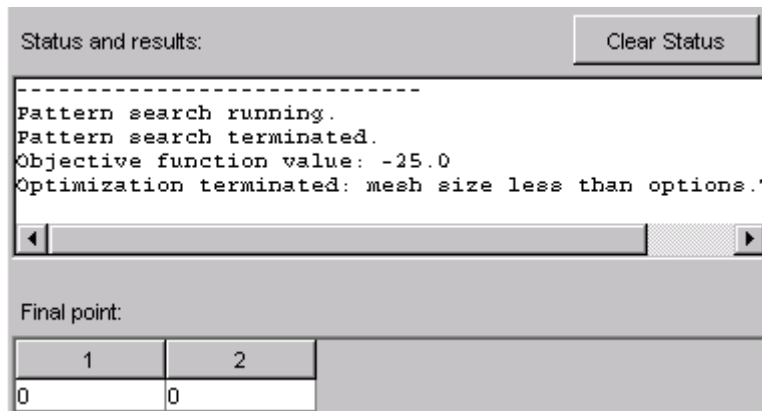
As soon as the search polls the mesh point (0, 6), at which the objective function value is less than at the initial point, it stops polling the current

mesh and sets the current point at the next iteration to (0, 6). Consequently, the search moves toward the local minimum at (0, 9) at the first iteration. You see this by looking at the first two lines of the command line display.

| Iter | f-count | MeshSize | f(x) | Method           |
|------|---------|----------|------|------------------|
| 0    | 1       | 1        | 0    | Start iterations |
| 1    | 3       | 2        | -7   | Successful Poll  |

Note that the pattern search performs only two evaluations of the objective function at the first iteration, increasing the total function count from 1 to 3.

Next, set **Complete poll** to 0n and click **Start**. The **Status and results** pane displays the following results.



This time, the pattern search finds the global minimum at (0, 0). The difference between this run and the previous one is that with **Complete poll** set to 0n, at the first iteration the pattern search polls all four mesh points.

$$f(0, 5) + (1, 0) = f(1, 5) = 0$$

$$f(0, 5) + (0, 1) = f(0, 6) = -6$$

$$f(0, 5) + (-1, 0) = f(-1, 5) = 0$$


$$f(0, 5) + (0, -1) = f(0, 4) = -9$$

Because the last mesh point has the lowest objective function value, the pattern search selects it as the current point at the next iteration. The first two lines of the command-line display show this.

| Iter | f-count | MeshSize | f(x) | Method           |
|------|---------|----------|------|------------------|
| 0    | 1       | 1        | 0    | Start iterations |
| 1    | 5       | 2        | -9   | Successful Poll  |

In this case, the objective function is evaluated four times at the first iteration. As a result, the pattern search moves toward the global minimum at (0, 0).

The following figure compares the sequence of points returned when **Complete poll** is set to `Off` with the sequence when **Complete poll** is `On`.



## Using a Search Method

In addition to polling the mesh points, the pattern search algorithm can perform an optional step at every iteration, called *search*. At each iteration,



the search step applies another optimization method to the current point. If this search does not improve the current point, the poll step is performed.

The following example illustrates the use of a search method on the problem described in “Example — A Linearly Constrained Problem” on page 6-6. To set up the example, enter the following commands at the MATLAB prompt to define the initial point and constraints.

```
x0 = [2 1 0 9 1 0];
Aineq = [-8 7 3 -4 9 0 ];
bineq = [7];
Aeq = [7 1 8 3 3 3; 5 0 5 1 5 8; 2 6 7 1 1 8; 1 0 0 0 0 0];
beq = [84 62 65 1];
```

Then enter the settings shown in the following figure in the Pattern Search Tool.

Objective function: @lincontest7

Start point: x0

Constraints:

Linear inequalities: A = Aineq b = bineq

Linear equalities: Aeq = Aeq beq = beq

Bounds: Lower = Upper =

Nonlinear constraint function:

Plots


Plot interval: 1

Best function value     Mesh size     Function count

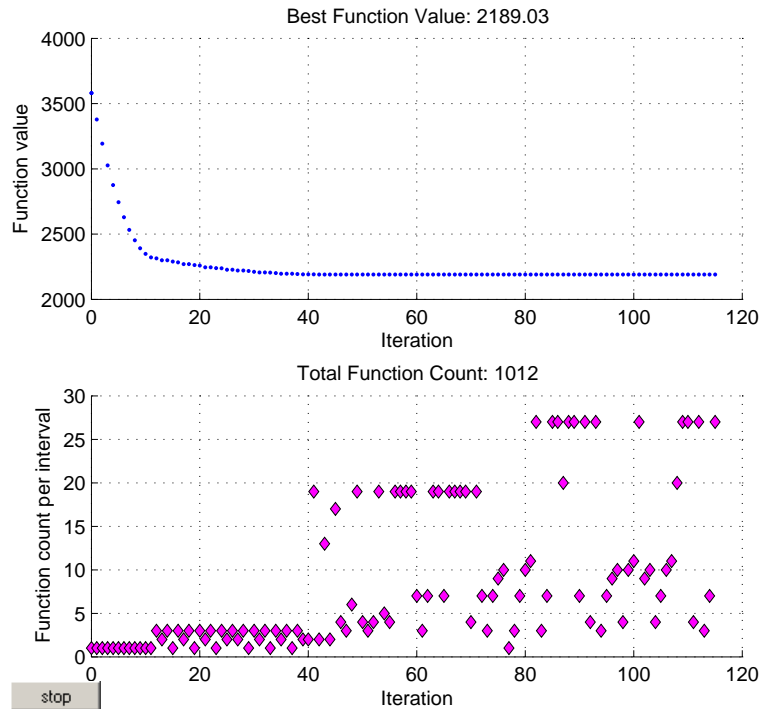
Best point     Max constraint

Custom function:

For comparison, click **Start** to run the example without a search method. This displays the plots shown in the following figure.



To see the effect of using a search method, select GPS Positive Basis Np1 in the **Search method** field in **Search** options. This sets the search method to be a pattern search using the pattern for GPS Positive basis Np1. Then click **Start** to run the genetic algorithm. This displays the following plots.



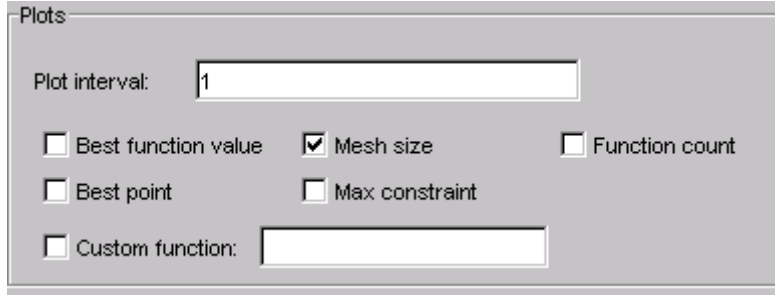
Note that using the search method reduces the total function count—the number of times the objective function was evaluated—by almost 50 percent, and reduces the number of iterations from 270 to 120.

## Mesh Expansion and Contraction


The **Expansion factor** and **Contraction factor** options, in **Mesh** options, control how much the mesh size is expanded or contracted at each iteration. With the default **Expansion factor** value of 2, the pattern search multiplies the mesh size by 2 after each successful poll. With the default **Contraction**

**factor** value of 0.5, the pattern search multiplies the mesh size by 0.5 after each unsuccessful poll.


You can view the expansion and contraction of the mesh size during the pattern search by selecting **Mesh size** in the **Plots** pane.



To also display the values of the mesh size and objective function at the command line, set **Level of display** to Iterative in the **Display to command window** options.




When you run the example described in “Example — A Linearly Constrained Problem” on page 6-6, the Pattern Search Tool displays the following plot.



To see the changes in mesh size more clearly, change the *y*-axis to logarithmic scaling as follows:

- 1** Select **Axes Properties** from the **Edit** menu in the plot window.
- 2** In the Properties Editor, select the **Y** tab.
- 3** Set **Scale** to **Log**.

Updating these settings in the MATLAB Property Editor will show the plot in the following figure.



The first 37 iterations result in successful polls, so the mesh sizes increase steadily during this time. You can see that the first unsuccessful poll occurs at iteration 38 by looking at the command-line display for that iteration.

|    |    |            |      |                 |
|----|----|------------|------|-----------------|
| 36 | 39 | 6.872e+010 | 3486 | Successful Poll |
| 37 | 40 | 1.374e+011 | 3486 | Successful Poll |
| 38 | 43 | 6.872e+010 | 3486 | Refine Mesh     |

Note that at iteration 37, which is successful, the mesh size doubles for the next iteration. But at iteration 38, which is unsuccessful, the mesh size is multiplied 0.5.

To see how **Expansion factor** and **Contraction factor** affect the pattern search, make the following changes:

- Set **Expansion factor** to 3.0.
- Set **Contraction factor** to 0.75.

Mesh

Initial size: 1

Max size: Inf

Accelerator: Off

Rotate: On

Scale: On

Expansion factor: 3.0

Contraction factor: 0.75

Then click **Start**. The **Status and results** pane shows that the final point is approximately the same as with the default settings of **Expansion factor** and **Contraction factor**, but that the pattern search takes longer to reach that point.

Status and results: Clear Status

```

-----
Pattern search running.
Pattern search terminated.
Objective function value: 2189.030445027741
Maximum number of iterations exceeded: increase option


```

Final point:

| 1     | 2        | 3       | 4        | 5        |
|-------|----------|---------|----------|----------|
| 1.001 | -2.30269 | 9.51315 | -0.04738 | -0.19772 |

The algorithm halts because it exceeds the maximum number of iterations, whose value you can set in the **Max iteration** field in the **Stopping criteria** options. The default value is 100 times the number of variables for the objective function, which is 6 in this example.

When you change the scaling of the  $y$ -axis to logarithmic, the mesh size plot appears as shown in the following figure.



Note that the mesh size increases faster with **Expansion factor** set to 3.0, as compared with the default value of 2.0, and decreases more slowly with **Contraction factor** set to 0.75, as compared with the default value of 0.5.

## Mesh Accelerator

The mesh accelerator can make a pattern search converge faster to the optimal point by reducing the number of iterations required to reach the mesh tolerance. When the mesh size is below a certain value, the pattern search contracts the mesh size by a factor smaller than the **Contraction factor** factor.

---

**Note** It is recommended to only use the mesh accelerator for problems in which the objective function is not too steep near the optimal point, or you might lose some accuracy. For differentiable problems, this means that the absolute value of the derivative is not too large near the solution.

---



To use the mesh accelerator, set **Accelerator** to On in the **Mesh** options. When you run the example described in “Example — A Linearly Constrained Problem” on page 6-6, the number of iterations required to reach the mesh tolerance is 246, as compared with 270 when **Accelerator** is set to Off.

You can see the effect of the mesh accelerator by setting **Level of display** to Iterative in **Display to command window**. Run the example with **Accelerator** set to On, and then run it again with **Accelerator** set to Off. The mesh sizes are the same until iteration 226, but differ at iteration 227. The MATLAB Command Window displays the following lines for iterations 226 and 227 with **Accelerator** set to Off.

| Iter | f-count | MeshSize   | f(x) | Method      |
|------|---------|------------|------|-------------|
| 226  | 1501    | 6.104e-005 | 2189 | Refine Mesh |
| 227  | 1516    | 3.052e-005 | 2189 | Refine Mesh |

Note that the mesh size is multiplied by 0.5, the default value of **Contraction factor**.

For comparison, the Command Window displays the following lines for the same iteration numbers with **Accelerator** set to On.

| Iter | f-count | MeshSize   | f(x) | Method      |
|------|---------|------------|------|-------------|
| 226  | 1501    | 6.104e-005 | 2189 | Refine Mesh |
| 227  | 1516    | 1.526e-005 | 2189 | Refine Mesh |

In this case the mesh size is multiplied by 0.25.

## Using Cache

Typically, at any given iteration of a pattern search, some of the mesh points might coincide with mesh points at previous iterations. By default, the pattern search recomputes the objective function at these mesh points even though it has already computed their values and found that they are not optimal. If computing the objective function takes a long time—say, several minutes—this can make the pattern search run significantly longer.

You can eliminate these redundant computations by using a cache, that is, by storing a history of the points that the pattern search has already visited. To do so, set **Cache** to On in **Cache** options. At each poll, the pattern search

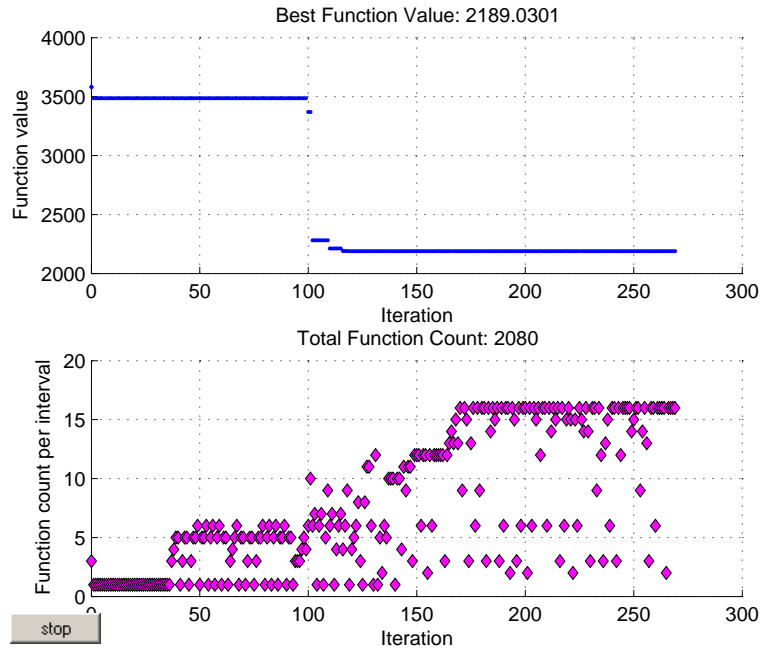
checks to see whether the current mesh point is within a specified tolerance, **Tolerance**, of a point in the cache. If so, the search does not compute the objective function for that point, but uses the cached function value and moves on to the next point.

---

**Note** When **Cache** is set to On, the pattern search might fail to identify a point in the current mesh that improves the objective function because it is within the specified tolerance of a point in the cache. As a result, the pattern search might run for more iterations with **Cache** set to On than with **Cache** set to Off. It is generally a good idea to keep the value of **Tolerance** very small, especially for highly nonlinear objective functions.


---

To illustrate this, select **Best function value** and **Function count** in the **Plots** pane and run the example described in “Example — A Linearly Constrained Problem” on page 6-6 with **Cache** set to Off. After the pattern search finishes, the plots appear as shown in the following figure.



Note that the total function count is 2080.

Now, set **Cache** to 0n and run the example again. This time, the plots appear as shown in the following figure.



This time, the total function count is reduced to 1973.

## Setting Tolerances for the Solver

Tolerance refers to how small a parameter, such as a mesh size, can become before the search is halted or changed in some way. You can specify the value of the following tolerances:

- **Mesh tolerance** — When the current mesh size is less than the value of **Mesh tolerance**, the algorithm halts.
- **X tolerance** — After a successful poll, if the distance from the previous best point to the current best point is less than the value of **X tolerance**, the algorithm halts.

- **Function tolerance** — After a successful poll, if the difference between the function value at the previous best point and function value at the current best point is less than the value of **Function tolerance**, the algorithm halts.
- **Nonlinear constraint tolerance** — The algorithm treats a point to be feasible if constraint violation is less than TolCon.
- **Bind tolerance** — Bind tolerance applies to constrained problems and specifies how close a point must get to the boundary of the feasible region before a linear constraint is considered to be active. When a linear constraint is active, the pattern search polls points in directions parallel to the linear constraint boundary as well as the mesh points.

Usually, you should set **Bind tolerance** to be at least as large as the maximum of **Mesh tolerance**, **X tolerance**, and **Function tolerance**.

### Example — Setting Bind Tolerance

The following example illustrates of how **Bind tolerance** affects a pattern search. The example finds the minimum of


$$f(x_1, x_2) = \sqrt{x_1^2 + x_2^2}$$

subject to the constraints

$$-11x_1 + 10x_2 \leq 10$$

$$10x_1 - 10x_2 \leq 10$$

Note that you can compute the objective function using the function `norm`. The feasible region for the problem lies between the two lines in the following figure.



### Running a Pattern Search with the Default Bind Tolerance

To run the example, enter `psearchtool` to open the Pattern Search Tool and enter the following information:


- Set **Objective function** to `@(x) norm(x)`.
- Set **Start point** to `[-1.001 -1.1]`.
- Select **Mesh size** in the Plots pane.
- Set **Level of display** to **Iterative** in the **Display to command window** options.

Then click **Start** to run the pattern search.

The display in the MATLAB Command Window shows that the first four polls are unsuccessful, because the mesh points do not lie in the feasible region.

| Iter | f-count | MeshSize | f(x)  | Method           |
|------|---------|----------|-------|------------------|
| 0    | 1       | 1        | 1.487 | Start iterations |
| 1    | 1       | 0.5      | 1.487 | Refine Mesh      |
| 2    | 1       | 0.25     | 1.487 | Refine Mesh      |
| 3    | 1       | 0.125    | 1.487 | Refine Mesh      |
| 4    | 1       | 0.0625   | 1.487 | Refine Mesh      |

The pattern search contracts the mesh at each iteration until one of the mesh points lies in the feasible region. The following figure shows a close-up of the initial point and mesh points at iteration 5.



The top mesh point, which is  $(-1.001, -1.0375)$ , has a smaller objective function value than the initial point, so the poll is successful.

Because the distance from the initial point to lower boundary line is less than the default value of **Bind tolerance**, which is 0.0001, the pattern search does not consider the linear constraint  $10x_1 - 10x_2 \leq 10$  to be active, so it does not search points in a direction parallel to the boundary line.

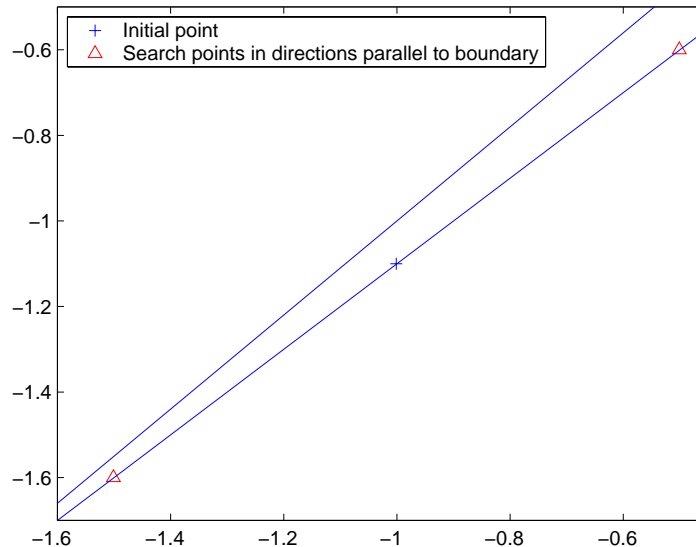
### Increasing the Value of Bind Tolerance

To see the effect of bind tolerance, change **Bind tolerance** to 0.01 and run the pattern search again.

This time, the display in the MATLAB Command Window shows that the first two iterations are successful.


| Iter | f-count | MeshSize | f(x)   | Method           |
|------|---------|----------|--------|------------------|
| 0    | 1       | 1        | 1.487  | Start iterations |
| 1    | 2       | 2        | 0.7817 | Successful Poll  |
| 2    | 3       | 4        | 0.6395 | Successful Poll  |

Because the distance from the initial point to the boundary is less than **Bind tolerance**, the second linear constraint is active. In this case, the pattern search polls points in directions parallel to the boundary line  $10x_1 - 10x_2 = 10$ , resulting in successful poll. The following figure shows the initial point with two addition search points in directions parallel to the boundary.





The following figure compares the sequences of points during the first 20 iterations of the pattern search for both settings of **Bind tolerance**.



Note that when **Bind tolerance** is set to .01, the points move toward the optimal point more quickly. The pattern search requires only 90 iterations. When **Bind tolerance** is set to .0001, the search requires 124 iterations. However, when the feasible region does not contain very acute angles, as it does in this example, increasing **Bind tolerance** can increase the number of iterations required, because the pattern search tends to poll more points.

## Constrained Minimization Using patternsearch

Suppose you want to minimize the simple objective function of two variables  $x_1$  and  $x_2$ ,

$$\min_x f(x) = (4 - 2.1 * x_1^2 - x_1^{\frac{4}{3}}) * x_1^2 + x_1 * x_2 + (-4 + 4 * x_2^2) * x_2^2$$

subject to the following nonlinear inequality constraints and bounds

$$x_1 \cdot x_2 + x_1 - x_2 + 1.5 \leq 0, \text{ (nonlinear constraint)}$$

$$10 - x_1 \cdot x_2 \leq 0, \text{ (nonlinear constraint)}$$

$$0 \leq x_1 \leq 1, \text{ (bound)}$$

$$0 \leq x_2 \leq 13. \text{ (bound)}$$

Begin by creating the objective and constraint functions. First, create an M-file named `simple_objective.m` as follows:

```
function y = simple_objective(x)
y = (4 - 2.1*x(1)^2 + x(1)^4/3)*x(1)^2 + x(1)*x(2) + (-4 + 4*x(2)^2)*x(2)^2;
```

The pattern search solver assumes the objective function will take one input  $x$  where  $x$  has as many elements as number of variables in the problem. The objective function computes the value of the function and returns that scalar value in its one return argument  $y$ .

Then create an M-file named `simple_constraint.m` containing the constraints:

```
function [c, ceq] = simple_constraint(x)
c = [1.5 + x(1)*x(2) + x(1) - x(2);
-x(1)*x(2) + 10];
ceq = [];
```

The pattern search solver assumes the constraint function will take one input  $x$ , where  $x$  has as many elements as the number of variables in the problem. The constraint function computes the values of all the inequality and equality constraints and returns two vectors,  $c$  and  $ceq$ , respectively.

Next, to minimize the objective function using the `patternsearch` function, you need to pass in a function handle to the objective function as well as specifying a start point as the second argument. Lower and upper bounds are provided as `LB` and `UB` respectively. In addition, you also need to pass a function handle to the nonlinear constraint function.

```
ObjectiveFunction = @simple_objective;
X0 = [0 0]; % Starting point
LB = [0 0]; % Lower bound
UB = [1 13]; % Upper bound
ConstraintFunction = @simple_constraint;
[x,fval] = patternsearch(ObjectiveFunction,X0,[],[],[],[],...
                        LB,UB,ConstraintFunction)
```

```
Optimization terminated: mesh size less than options.TolMesh
and constraint violation is less than options.TolCon.
```

```
x =
```

```
    0.8122    12.3122
```

```
fval =
```

```
    9.1324e+004
```

Next, plot the results. Create an options structure using `psoptimset` that selects two plot functions. The first plot function `psplotbestf` plots the best objective function value at every iteration. The second plot function `psplotmaxconstr` plots the maximum constraint violation at every iteration.

---

**Note** You can also visualize the progress of the algorithm by displaying information to the Command Window using the `'Display'` option.

---

```
options = psoptimset('PlotFcns',{@psplotbestf,@psplotmaxconstr},'Display','iter');
[x,fval] = patternsearch(ObjectiveFunction,X0,[],[],[],[],LB,UB,ConstraintFunction,options)
```

| Iter | f-count | f(x)    | max<br>constraint | MeshSize | Method           |
|------|---------|---------|-------------------|----------|------------------|
| 0    | 1       | 0       | 10                | 0.8919   |                  |
| 1    | 5       | 113580  | 0                 | 0.001    | Increase penalty |
| 2    | 24      | 91324.4 | 0                 | 1e-005   | Increase penalty |
| 3    | 48      | 91324   | 0                 | 1e-007   | Increase penalty |

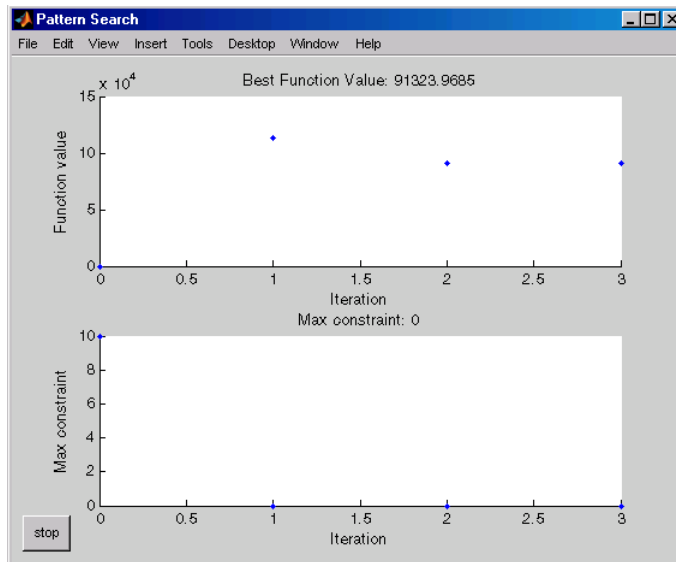
Optimization terminated: mesh size less than options.TolMesh  
and constraint violation is less than options.TolCon.

x =

0.8122 12.3122

fval =

9.1324e+004



**Best Objective Function Value and Maximum Constraint Violation at Each Iteration**

## Parameterizing Functions Called by patternsearch

Sometimes you might want to write functions that are called by `patternsearch` that have additional parameters to the independent variable. For example, suppose you want to minimize the following function:

$$f(x) = (a - bx_1^2 + x_1^{4/3})x_1^2 + x_1x_2 + (-c + cx_2^2)x_2^2$$

for different values of  $a$ ,  $b$ , and  $c$ . Because `patternsearch` accepts objective functions that depend only on  $x$ , you must provide the additional parameters  $a$ ,  $b$ , and  $c$  to the function before calling `patternsearch`. The following sections describe two ways to do this:

- “Parameterizing Functions Using Anonymous Functions with `patternsearch`” on page 6-53
- “Parameterizing a Function Using a Nested Function with `patternsearch`” on page 6-55

The examples in these sections show how to parameterize the objective function, but you can use the same methods to parameterize any user-defined functions called by `patternsearch`, for example using a custom search method.

### Parameterizing Functions Using Anonymous Functions with `patternsearch`

To parameterize your function, first write an M-file containing the following code:

```
function y = parameterfun(x,a,b,c)
y = (a - b*x(1)^2 + x(1)^4/3)*x(1)^2 + x(1)*x(2) + ...
    (-c + c*x(2)^2)*x(2)^2;
```

Save the M- file as `parameterfun.m` in a directory on the MATLAB path.


Now, suppose you want to minimize the function for the parameter values  $a = 4$ ,  $b = 2.1$ , and  $c = 4$ . To do so, define a function handle to an anonymous function by entering the following commands at the MATLAB prompt:

```
a = 4; b = 2.1; c = 4;    % Define parameter values
objfun = @(x) parameterfun(x,a,b,c);
x0 = [0.5 0.5];
```

If you are using the Pattern Search Tool,

- Set **Objective function** to objfun.
- Set **Start point** to x0.


The following figure shows these settings in the Pattern Search Tool.



Objective function:

Start point:

Next, click **Start** to run the optimization. The **Status and results** pane displays the final answer.



Run solver

Current iteration:

Status and results:

-----  
Pattern search running.  
Pattern search terminated.  
Objective function value: -1.0316284534886349  
Optimization terminated: mesh size less than options.  
-----

Final point:

| 1        | 2       |
|----------|---------|
| -0.08984 | 0.71266 |

If you subsequently decide to change the values of  $a$ ,  $b$ , and  $c$ , you must recreate the anonymous function. For example,

```
a = 3.6; b = 2.4; c = 5;    % Define parameter values
objfun = @(x) parameterfun(x,a,b,c);
```

## Parameterizing a Function Using a Nested Function with patternsearch

As an alternative to parameterizing the objective function as an anonymous function, you can write a single M-file that

- Accepts  $a$ ,  $b$ ,  $c$ , and  $x_0$  as inputs.
- Contains the objective function as a nested function.
- Calls `patternsearch`.

The following shows the code for the M-file.

```
function [x fval] = runps(a,b,c,x0)
[x, fval] = patternsearch(@nestedfun,x0);
% Nested function that computes the objective function
function y = nestedfun(x)
    y = (a - b*x(1)^2 + x(1)^4/3)*x(1)^2 + x(1)*x(2) + ...
        (-c + c*x(2)^2)*x(2)^2;
end
end
```

Note that the objective function is computed in the nested function `nestedfun`, which has access to the variables `a`, `b`, and `c`. To run the optimization, enter

```
[x fval] = runps(a,b,c,x0)
```

This returns

```
Optimization terminated: mesh size less than options.TolMesh.
```

```
x =
```

```
   -0.0898    0.7127
```

```
fval =
```

```
   -1.0316
```



# Using the Genetic Algorithm

---

Overview of the Genetic Algorithm Tool (p. 7-2)

Provides an overview of the Genetic Algorithm Tool.

Using the Genetic Algorithm from the Command Line (p. 7-24)

Describes how to use the genetic algorithm at the command line.

Genetic Algorithm Examples (p. 7-34)

Explains how to set options for the genetic algorithm.

Parameterizing Functions Called by ga (p. 7-67)

Explains how to write functions with additional parameters.

## Overview of the Genetic Algorithm Tool

The section provides an overview of the Genetic Algorithm Tool. This section covers the following topics:

- “Opening the Genetic Algorithm Tool” on page 7-2
- “Defining an Unconstrained Problem in the Genetic Algorithm Tool” on page 7-4
- “Running the Genetic Algorithm” on page 7-4
- “Pausing and Stopping the Algorithm” on page 7-6
- “Displaying Plots” on page 7-8
- “Example — Creating a Custom Plot Function” on page 7-9
- “Reproducing Your Results” on page 7-12
- “Setting Options in the Genetic Algorithm Tool” on page 7-14
- “Importing and Exporting Options and Problems” on page 7-15
- “Example — Resuming the Genetic Algorithm from the Final Population” on page 7-19


### Opening the Genetic Algorithm Tool

To open the tool, enter

```
gatool
```

at the MATLAB prompt.

This opens the Genetic Algorithm Tool, as shown in the following figure.



### Defining an Unconstrained Problem in the Genetic Algorithm Tool

You can define the problem you want to solve in the following two fields:

- **Fitness function** — The function you want to minimize. Enter a handle to an M-file function that computes the fitness function. Chapter 2, “Writing M-Files for Functions You Want to Optimize” describes how to write the M-file.
- **Number of variables** — The number of independent variables for the fitness function.

The following figure shows these fields for the example described in “Example: Rastrigin’s Function” on page 4-7.




The image shows a screenshot of a software interface with two input fields. The first field is labeled "Fitness function:" and contains the text "@rastriginsfcn". The second field is labeled "Number of variables:" and contains the number "2".

### Running the Genetic Algorithm

To run the genetic algorithm, click **Start** in the **Run solver** pane. When you do so,

- The **Current generation** field displays the number of the current generation.
- The **Status and results** pane displays the message GA running.

The following figure shows the **Current generation** field and **Status and results** pane.



When the algorithm terminates, the **Status and results** pane displays

- The message GA terminated.
- The fitness function value of the best individual in the final generation
- The reason the algorithm terminated
- The coordinates of the final point

You can change many of the settings in the Genetic Algorithm Tool while the algorithm is running. Your changes are applied at the next generation. Until your changes are applied, which occurs at the start of the next generation, the **Status and results** pane displays the message Changes pending. At the start of the next generation, the pane displays the message Changes applied.

### Pausing and Stopping the Algorithm

While the genetic algorithm is running, you can

- Click **Pause** to temporarily suspend the algorithm. To resume the algorithm using the current population at the time you paused, click **Resume**.
- Click **Stop** to stop the algorithm. The **Status and results** pane displays the fitness function value of the best point in the current generation at the moment you clicked **Stop**.

---

**Note** If you click **Stop** and then run the genetic algorithm again by clicking **Start**, the algorithm begins with a new random initial population or with the population you specify in the **Initial population** field. If you want to restart the algorithm where it left off, use the **Pause** and **Resume** buttons.

---

“Example — Resuming the Genetic Algorithm from the Final Population” on page 7-19 explains what to do if you click **Stop** and later decide to resume the genetic algorithm from the final population of the last run.

### Setting Stopping Criteria

The genetic algorithm uses several criteria, listed in the **Stopping criteria** options, to decide when to stop, in case you do not stop it manually by clicking **Stop**. The algorithm stops if any one of the following conditions occur:

- **Generations** — The algorithm reaches the specified number of generations.
- **Time limit** — The algorithm runs for the specified amount of time in seconds.
- **Fitness limit** — The best fitness value in the current generation is less than or equal to the specified value.
- **Stall generations** — The algorithm computes the specified number of generations with no improvement in the fitness function.
- **Stall time limit** — The algorithm runs for the specified amount of time in seconds with no improvement in the fitness function.


- **Function tolerance** — The algorithm runs until the cumulative change in the fitness function value over **Stall generations** is less than **Function Tolerance**.

The **Nonlinear constraint tolerance** is not used as stopping criterion. It is used to determine the feasibility with respect to nonlinear constraints.

If you want the genetic algorithm to continue running until you click **Pause** or **Stop**, you should change the default values of the following five options as follows:

- Set **Generations** to Inf.
- Set **Time limit** to Inf.
- Set **Fitness limit** to -Inf.
- Set **Stall generations** to Inf.
- Set **Stall time limit** to Inf.

The following figure shows these settings.



| Stopping criteria               |        |
|---------------------------------|--------|
| Generations:                    | Inf    |
| Time limit:                     | Inf    |
| Fitness limit:                  | -Inf   |
| Stall generations:              | Inf    |
| Stall time limit:               | Inf    |
| Function tolerance:             | 1e-006 |
| Nonlinear constraint tolerance: | 1e-006 |


---

**Note** Do not use these settings when calling the genetic algorithm function `ga` at the command line. The function will not terminate until you press **Ctrl + C**. Instead, set **Generations** or **Time limit** to a finite number.

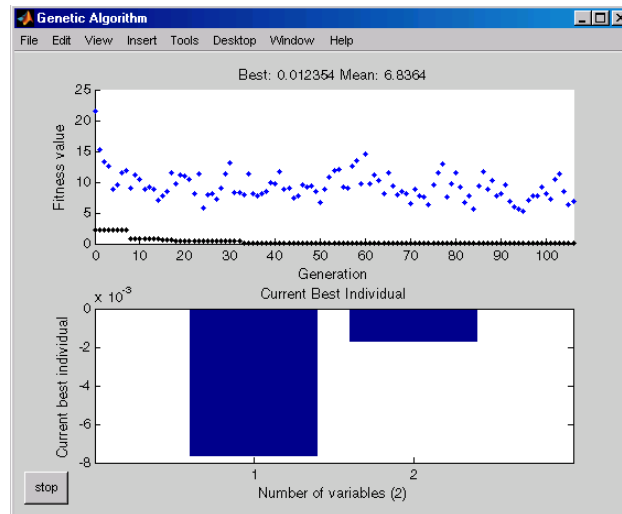
---

## Displaying Plots

The **Plots** pane, shown in the following figure, enables you to display various plots of the results of the genetic algorithm.



Select the check boxes next to the plots you want to display. For example, if you select **Best fitness** and **Best individual**, and run the example described in “Example: Rastrigin’s Function” on page 4-7, the tool displays plots similar to those shown in the following figure.



The upper plot displays the best and mean fitness values in each generation. The lower plot displays the coordinates of the point with the best fitness value in the current generation.



---

**Note** When you display more than one plot, clicking on any plot while the genetic algorithm is running or after the solver has completed opens a larger version of the plot in a separate window.

---

“Plot Options” on page 9-24 describes the types of plots you can create.

## Example – Creating a Custom Plot Function

If none of the plot functions that come with the toolbox is suitable for the output you want to plot, you can write your own custom plot function, which the genetic algorithm calls at each generation to create the plot. This example shows how to create a plot function that displays the change in the best fitness value from the previous generation to the current generation.

This section covers the following topics:

- “Creating the Custom Plot Function” on page 7-9
- “Using the Plot Function” on page 7-11
- “How the Plot Function Works” on page 7-11

## Creating the Custom Plot Function

To create the plot function for this example, copy and paste the following code into a new M-file in the MATLAB Editor.

```
function state = gaplotchange(options, state, flag)
% GAPLOTCHANGE Plots the logarithmic change in the best score from the
% previous generation.
%
persistent last_best % Best score in the previous generation

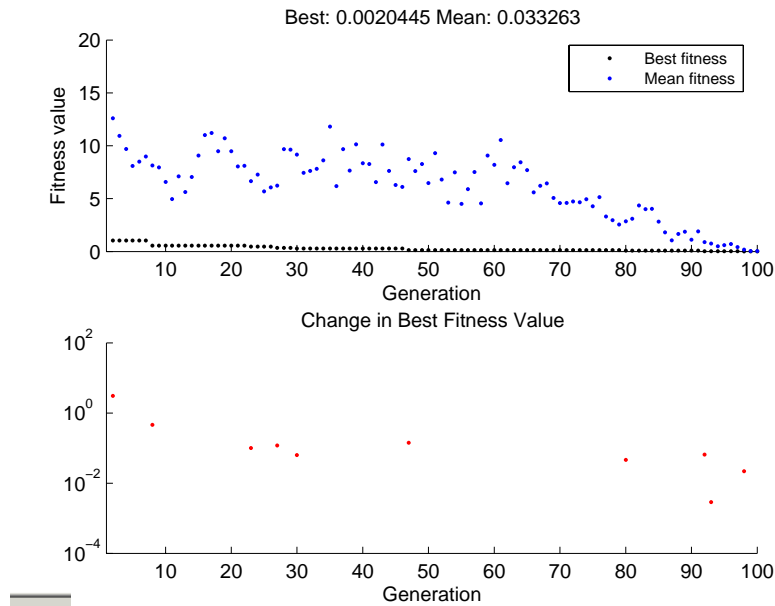
if(strcmp(flag,'init')) % Set up the plot
    set(gca,'xlim',[1,options.Generations],'yscale','log');
    hold on;
    xlabel Generation
    title('Change in Best Fitness Value')
end
```

```
best = min(state.Score); % Best score in the current generation
if state.Generation == 0 % Set last_best to best.
    last_best = best;
else
    change = last_best - best; % Change in best score
    last_best=best;
    plot(state.Generation, change, '.r');
    title(['Change in Best Fitness Value'])
end
```

Then save the M-file as `gaplotchange.m` in a directory on the MATLAB path.

## Using the Plot Function

To use the custom plot function, select **Custom** in the **Plots** pane and enter `@gaplotchange` in the field to the right. To compare the custom plot with the best fitness value plot, also select **Best fitness**. Now, if you run the example described in “Example: Rastrigin’s Function” on page 4-7, the tool displays plots similar to those shown in the following figure.



Note that because the scale of the  $y$ -axis in the lower custom plot is logarithmic, the plot only shows changes that are greater than 0. The logarithmic scale enables you to see small changes in the fitness function that the upper plot does not reveal.

## How the Plot Function Works

The plot function uses information contained in the following structures, which the genetic algorithm passes to the function as input arguments:

- `options` — The current options settings
- `state` — Information about the current generation

- `flag` — String indicating the current status of the algorithm

The most important lines of the plot function are the following:

- `persistent last_best`

Creates the persistent variable `last_best`—the best score in the previous generation. Persistent variables are preserved over multiple calls to the plot function.

- `set(gca, 'xlim', [1, options.Generations], 'Yscale', 'log');`

Sets up the plot before the algorithm starts. `options.Generations` is the maximum number of generations.

- `best = min(state.Score)`

The field `state.Score` contains the scores of all individuals in the current population. The variable `best` is the minimum score. For a complete description of the fields of the structure `state`, see “Structure of the Plot Functions” on page 9-25.

- `change = last_best - best`

The variable `change` is the best score at the previous generation minus the best score in the current generation.

- `plot(state.Generation, change, 'r')`

Plots the change at the current generation, whose number is contained in `state.Generation`.

The code for `gplotchange` contains many of the same elements as the code for `gplotbestf`, the function that creates the best fitness plot.

## Reproducing Your Results

To reproduce the results of the last run of the genetic algorithm, select the **Use random states from previous run** check box. This resets the states of the random number generators used by the algorithm to their previous values. If you do not change any other settings in the Genetic Algorithm Tool, the next time you run the genetic algorithm, it returns the same results as the previous run.

Normally, you should leave **Use random states from previous run** unselected to get the benefit of randomness in the genetic algorithm. Select the **Use random states from previous run** check box if you want to analyze the results of that particular run or show the exact results to others. After the algorithm has run, you can clear your results using the **Clear Status** button in the **Run solver** settings.

## Setting Options in the Genetic Algorithm Tool

You can set options for the genetic algorithm in the **Options** pane, shown in the figure below.

The image shows a software interface for setting genetic algorithm options. The window is titled "Options:" and has a ">>" button in the top right corner. The options are organized into several sections, each with a plus sign icon in a square:

- Population**:
  - Population type: Double Vector (dropdown menu)
  - Population size: 20 (text input)
  - Creation function: Uniform (dropdown menu)
  - Initial population: [] (text input)
  - Initial scores: [] (text input)
  - Initial range: [0 ; 1 ] (text input)
- Fitness scaling**
- Selection**
- Reproduction**
- Mutation**
- Crossover**
- Migration**
- Algorithm settings**
- Hybrid function**
- Stopping criteria**
- Output function**
- Display to command window**
- Vectorize**

“Genetic Algorithm Examples” on page 7-34 describes how options settings affect the performance of the genetic algorithm. For a detailed description of all the available options, see “Genetic Algorithm Options” on page 9-23.

### Setting Options as Variables in the MATLAB Workspace

You can set numerical options either directly, by typing their values in the corresponding edit box, or by entering the name of a variable in the MATLAB workspace that contains the option values. For example, you can set the **Population size** to 50 in either of the following ways:

- Enter 50 in the **Population size** field.
- Enter

```
popsize = 50
```

at the MATLAB prompt and then enter popsize in the **Population size** field.

For options whose values are large matrices or vectors, it is often more convenient to define their values as variables in the MATLAB workspace. This way, it is easy to change the entries of the matrix or vector if necessary.

### Importing and Exporting Options and Problems

You can export options and problem structures from the Genetic Algorithm Tool to the MATLAB workspace, and then import them back into the tool at a later time. This enables you to save the current settings for a problem and restore them later. You can also export the options structure and use it with the genetic algorithm function `ga` at the command line.

You can import and export the following information:

- The problem definition, including **Fitness function**, **Number of variables**, **Linear inequalities**, **Linear equalities**, **Bounds**, and a **Nonlinear constraint function**
- The currently specified options
- The results of the algorithm

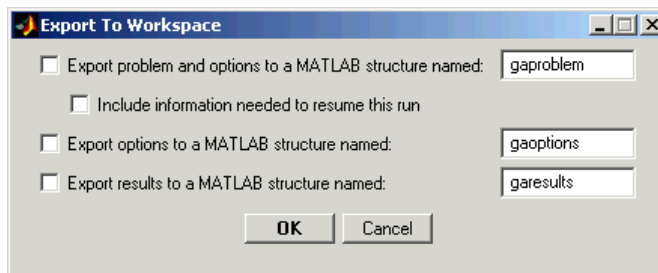
The following sections explain how to import and export this information:

- “Exporting Options and Problems” on page 7-16
- “Example — Running ga on an Exported Problem” on page 7-17
- “Importing Options” on page 7-18
- “Importing Problems” on page 7-19
- “Resetting the Problem Fields” on page 7-19

### Exporting Options and Problems

You can export options and problems to the MATLAB workspace so that you can use them at a future time in the Genetic Algorithm Tool. You can also apply the function `ga` using these options or problems at the command line—see “Using Options and Problems from the Genetic Algorithm Tool” on page 7-28.

To export options or problems, click the **Export to Workspace** button or select **Export to Workspace** from the **File** menu. This opens the dialog box shown in the following figure.



The dialog box provides the following options:

- To save both the problem definition and the current options settings, select **Export problem and options to a MATLAB structure named** and enter a name for the structure. Clicking **OK** saves this information to a structure in the MATLAB workspace. If you later import this structure into the Genetic Algorithm Tool, the settings for **Fitness function**, **Number of variables**, and all options settings are restored to the values they had when you exported the structure.



---

**Note** If you select **Use random states from previous run** in the **Run solver** pane before exporting a problem, the Genetic Algorithm Tool also saves the states of `rand` and `randn` at the beginning of the last run when you export. Then, when you import the problem and run the genetic algorithm with **Use random states from previous run** selected, the results of the run just before you exported the problem are reproduced exactly.

---

- If you want the genetic algorithm to resume from the final population of the last run before you exported the problem, select **Include information needed to resume this run**. Then, when you import the problem structure and click **Start**, the algorithm resumes from the final population of the previous run.

To restore the genetic algorithm's default behavior of generating a random initial population, delete the population in the **Initial population** field and replace it with empty brackets, `[]`.

---

**Note** If you select **Include information needed to resume this run**, then selecting **Use random states from previous run** has no effect on the initial population created when you import the problem and run the genetic algorithm on it. The latter option is only intended to reproduce results from the beginning of a new run, not from a resumed run.

---

- To save only the options, select **Export options to a MATLAB structure named** and enter a name for the options structure.
- To save the results of the last run of the algorithm, select **Export results to a MATLAB structure named** and enter a name for the results structure.

### **Example – Running ga on an Exported Problem**

To export the problem described in “Example: Rastrigin’s Function” on page 4-7 and run the genetic algorithm function `ga` on it at the command line, do the following steps:

- 1 Click **Export to Workspace**.

- 2 In the Export to Workspace dialog box, enter a name for the problem structure, such as `my_problem`, in the **Export problems and options to a MATLAB structure named** field.
- 3 At the MATLAB prompt, call the function `ga` with `my_problem` as the input argument:

```
[x fval] = ga(my_problem)
```

This returns

```
x =  
  
    0.0027   -0.0052  
  
fval =  
  
    0.0068
```

See “Using the Genetic Algorithm from the Command Line” on page 7-24 for form information.

### Importing Options

To import an options structure from the MATLAB workspace, select **Import Options** from the **File** menu. This opens a dialog box that displays a list of the genetic algorithm options structures in the MATLAB workspace. When you select an options structure and click **Import**, the options fields in the Genetic Algorithm Tool are updated to display the values of the imported options.

You can create an options structure in either of the following ways:

- Calling `gaoptimset` with options as an output argument
- By exporting the current options from the Export to Workspace dialog box in the Genetic Algorithm Tool

## Importing Problems

To import a problem that you previously exported from the Genetic Algorithm Tool, select **Import Problem** from the **File** menu. This opens the dialog box that displays a list of the genetic algorithm problem structures in the MATLAB workspace. When you select a problem structure and click **OK**, the following fields are updated in the Genetic Algorithm Tool:

- **Fitness function**
- **Number of variables**
- **Linear inequalities**
- **Linear equalities**
- **Bounds**
- **Nonlinear constraint function**
- The fields specified in the **Options** panel

## Resetting the Problem Fields


To reset or clear a problem you previously ran in the Genetic Algorithm Tool, select **New Problem** from the **File** menu. This sets all the fields in the Genetic Algorithm Tool to the default options.

## Example – Resuming the Genetic Algorithm from the Final Population

The following example shows how export a problem so that when you import it and click **Start**, the genetic algorithm resumes from the final population saved with the exported problem. To run the example, enter the following information in the Genetic Algorithm Tool:

- Set **Fitness function** to `@ackleyfcn`, which computes Ackley's function, a test function provided with the toolbox.
- Set **Number of variables** to 10.
- Select **Best fitness** in the **Plots** pane.
- Click **Start**.

This displays the following plot.



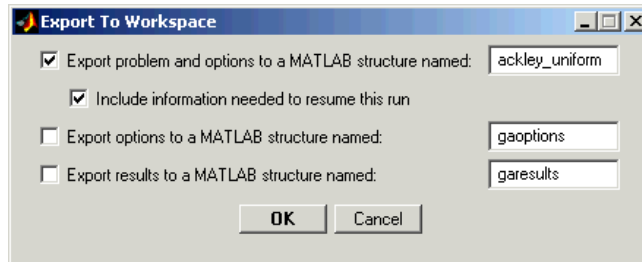
Suppose you want to experiment by running the genetic algorithm with other options settings, and then later restart this run from its final population with its current options settings. You can do this using the following steps:

**1** Click **Export to Workspace**.

**2** In the dialog box that appears,

- Select **Export problem and options to a MATLAB structure named**.
- Enter a name for the problem and options, such as `ackley_uniform`, in the text field.
- Select **Include information needed to resume this run**.

The dialog box should now appear as in the following figure.



### 3 Click **OK**.

This exports the problem and options to a structure in the MATLAB workspace. You can view the structure in the MATLAB Command Window by entering

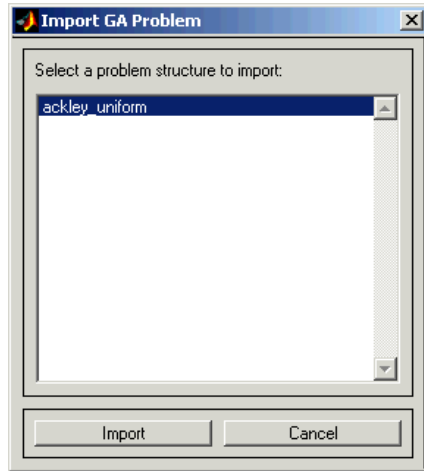
```
ackley_uniform

ackley_uniform =

    fitnessfcn: @ackleyfcn
           nvars: 10
           LB: []
           UB: []
           Aineq: []
           bineq: []
           Aeq: []
           beq: []
           nonlcon: []
           options: [1x1 struct]
```


After running the genetic algorithm with different options settings or even a different fitness function, you can restore the problem as follows:

- 1 Select **Import Problem** from the **File** menu. This opens the dialog box shown in the following figure.



- 2 Select `ackley_uniform`.
- 3 Click **Import**.

This sets the **Initial population** and **Initial scores** fields in the **Population** panel to the final population of the run before you exported the problem. All other options are restored to their setting during that run. When you click **Start**, the genetic algorithm resumes from the saved final population. The following figure shows the best fitness plots from the original run and the restarted run.




---

**Note** If, after running the genetic algorithm with the imported problem, you want to restore the genetic algorithm's default behavior of generating a random initial population, delete the population in the **Initial population** field.

---

## Generating an M-File

To create an M-file that runs the genetic algorithm, using the fitness function and options you specify in the Genetic Algorithm Tool, select **Generate M-File** from the **File** menu and save the M-file in a directory on the MATLAB path. Calling this M-file at the command line returns the same results as the Genetic Algorithm Tool, using the fitness function and options settings that were in place when you generated the M-file.

## Using the Genetic Algorithm from the Command Line

As an alternative to using the Genetic Algorithm Tool, you can run the genetic algorithm function `ga` from the command line. This section explains how to do so and covers the following topics.

- “Running `ga` with the Default Options” on page 7-24
- “Setting Options for `ga` at the Command Line” on page 7-26
- “Using Options and Problems from the Genetic Algorithm Tool” on page 7-28
- “Reproducing Your Results” on page 7-29
- “Resuming `ga` from the Final Population of a Previous Run” on page 7-30
- “Running `ga` from an M-File” on page 7-31

### Running `ga` with the Default Options

To run the genetic algorithm with the default options, call `ga` with the syntax

```
[x fval] = ga(@fitnessfun, nvars)
```

The input arguments to `ga` are

- `@fitnessfun` — A function handle to the M-file that computes the fitness function. Chapter 2, “Writing M-Files for Functions You Want to Optimize” explains how to write this M-file.
- `nvars` — The number of independent variables for the fitness function.

The output arguments are

- `x` — The final point
- `fval` — The value of the fitness function at `x`

For a description of additional input and output arguments, see the reference page for `ga`.



You can run the example described in “Example: Rastrigin’s Function” on page 4-7 from the command line by entering

```
[x fval] = ga(@rastriginsfcn, 2)
```

This returns

```
x =  
  
    0.0027   -0.0052  
  
fval =  
  
    0.0068
```

### Additional Output Arguments

To get more information about the performance of the genetic algorithm, you can call `ga` with the syntax

```
[x fval exitflag output population scores] = ga(@fitnessfcn, nvars)
```

Besides `x` and `fval`, this function returns the following additional output arguments:

- `exitflag` — Integer value corresponding to the reason the algorithm terminated
- `output` — Structure containing information about the performance of the algorithm at each generation
- `population` — Final population
- `scores` — Final scores

See the `ga` reference page for more information about these arguments.

## Setting Options for ga at the Command Line

You can specify any of the options that are available in the Genetic Algorithm Tool by passing an options structure as an input argument to ga using the syntax

```
[x fval] = ga(@fitnessfun, nvars, [],[],[],[],[],[],[],options)
```

This syntax does not specify any linear equality, linear inequality, or nonlinear constraints.

You create the options structure using the function gaoptimset.

```
options = gaoptimset
```

This returns the structure options with the default values for its fields.

```
options =  
  
    PopulationType: 'doubleVector'  
    PopInitRange: [2x1 double]  
    PopulationSize: 20  
    EliteCount: 2  
    CrossoverFraction: 0.8000  
    MigrationDirection: 'forward'  
    MigrationInterval: 20  
    MigrationFraction: 0.2000  
    Generations: 100  
    TimeLimit: Inf  
    FitnessLimit: -Inf  
    StallGenLimit: 50  
    StallTimeLimit: 20  
    TolFun: 1.0000e-006  
    TolCon: 1.0000e-006  
    InitialPopulation: []  
    InitialScores: []  
    InitialPenalty: 10  
    PenaltyFactor: 100  
    PlotInterval: 1  
    CreationFcn: @gacreationuniform  
    FitnessScalingFcn: @fitscalingrank
```

```

SelectionFcn: @selectionstochunif
CrossoverFcn: @crossoverscattered
MutationFcn: @mutationgaussian
HybridFcn: []
Display: 'final'
PlotFcns: []
OutputFcns: []
Vectorized: 'off'

```

The function `ga` uses these default values if you do not pass in options as an input argument.

The value of each option is stored in a field of the options structure, such as `options.PopulationSize`. You can display any of these values by entering options followed by the name of the field. For example, to display the size of the population for the genetic algorithm, enter

```
options.PopulationSize
```

```
ans =
```

```
20
```

To create an options structure with a field value that is different from the default — for example to set `PopulationSize` to 100 instead of its default value 20 — enter

```
options = gaoptimset('PopulationSize', 100)
```

This creates the options structure with all values set to their defaults except for `PopulationSize`, which is set to 100.

If you now enter,

```
ga(@fitnessfun,nvars,[],[],[],[],[],[],[],options)
```

`ga` runs the genetic algorithm with a population size of 100.

If you subsequently decide to change another field in the options structure, such as setting `PlotFcns` to `@gaplotbestf`, which plots the best fitness function value at each generation, call `gaoptimset` with the syntax

```
options = gaoptimset(options, 'PlotFcns', @plotbestf)
```

This preserves the current values of all fields of options except for `PlotFcns`, which is changed to `@plotbestf`. Note that if you omit the input argument `options`, `gaoptimset` resets `PopulationSize` to its default value 20.

You can also set both `PopulationSize` and `PlotFcns` with the single command

```
options = gaoptimset('PopulationSize',100,'PlotFcns',@plotbestf)
```

## Using Options and Problems from the Genetic Algorithm Tool

As an alternative to creating an options structure using `gaoptimset`, you can set the values of options in the Genetic Algorithm Tool and then export the options to a structure in the MATLAB workspace, as described in “Exporting Options and Problems” on page 7-16. If you export the default options in the Genetic Algorithm Tool, the resulting structure `options` has the same settings as the default structure returned by the command

```
options = gaoptimset
```

If you export a problem structure, `ga_problem`, from the Genetic Algorithm Tool, you can apply `ga` to it using the syntax

```
[x fval] = ga(ga_problem)
```

The problem structure contains the following fields:

- `fitnessfcn` — Fitness function
- `nvars` — Number of variables for the problem
- `Aineq` — Matrix for inequality constraints
- `Bineq` — Vector for inequality constraints
- `Aeq` — Matrix for equality constraints

- Beq — Vector for equality constraints
- LB — Lower bound on x
- UB — Upper bound on x
- nonlcon — Nonlinear constraint function
- options — Options structure

## Reproducing Your Results

Because the genetic algorithm is stochastic—that is, it makes random choices—you get slightly different results each time you run the genetic algorithm. The algorithm uses the MATLAB uniform and normal random number generators, `rand` and `randn`, to make random choices at each iteration. Each time `ga` calls `rand` and `randn`, their states are changed, so that the next time they are called, they return different random numbers. This is why the output of `ga` differs each time you run it.

If you need to reproduce your results exactly, you can call `ga` with an output argument that contains the current states of `rand` and `randn` and then reset the states to these values before running `ga` again. For example, to reproduce the output of `ga` applied to Rastrigin's function, call `ga` with the syntax

```
[x fval exitflag output] = ga(@rastriginsfcn, 2);
```

Suppose the results are

```
x =  
    0.0027   -0.0052  
fval =  
    0.0068
```

The states of `rand` and `randn` are stored in the first two fields of output.

```
output =  
  
    randstate: [35x1 double]  
    randnstate: [2x1 double]  
    generations: 100  
    funccount: 2000  
    message: [1x64 char]
```

Then, reset the states, by entering

```
rand('twister', output.randstate);  
randn('state', output.randnstate);
```

If you now run `ga` a second time, you get the same results.

---

**Note** If you do not need to reproduce your results, it is better not to set the states of `rand` and `randn`, so that you get the benefit of the randomness in the genetic algorithm.

---

### Resuming `ga` from the Final Population of a Previous Run

By default, `ga` creates a new initial population each time you run it. However, you might get better results by using the final population from a previous run as the initial population for a new run. To do so, you must have saved the final population from the previous run by calling `ga` with the syntax

```
[x, fval, exitflag, output, final_pop] = ga(@fitnessfcn, nvars);
```

The last output argument is the final population. To run `ga` using `final_pop` as the initial population, enter

```
options = gaoptimset('InitialPop', final_pop);  
[x, fval, exitflag, output, final_pop2] = ...  
    ga(@fitnessfcn, nvars,[],[],[],[],[],[],[],[],options);
```

You can then use `final_pop2`, the final population from the second run, as the initial population for a third run.

## Running ga from an M-File


The command-line interface enables you to run the genetic algorithm many times, with different options settings, using an M-file. For example, you can run the genetic algorithm with different settings for **Crossover fraction** to see which one gives the best results. The following code runs the function `ga` 21 times, varying `options.CrossoverFraction` from 0 to 1 in increments of 0.05, and records the results.

```
options = gaoptimset('Generations',300);
rand('twister', 71); % These two commands are only included to
randn('state', 59); % make the results reproducible.
record=[];
for n=0:.05:1
    options = gaoptimset(options,'CrossoverFraction', n);
    [x fval]=ga(@rastriginsfcn, 10,[],[],[],[],[],[],[],options);
    record = [record; fval];
end
```

You can plot the values of `fval` against the crossover fraction with the following commands:

```
plot(0:.05:1, record);
xlabel('Crossover Fraction');
ylabel('fval')
```


The following plot appears.



The plot indicates that you get the best results by setting options.CrossoverFraction to a value somewhere between 0.6 and 0.95.



You can get a smoother plot of `fval` as a function of the crossover fraction by running `ga` 20 times and averaging the values of `fval` for each crossover fraction. The following figure shows the resulting plot.



The plot narrows the range of best choices for `options.CrossoverFraction` to values between 0.7 and 0.9.

## Genetic Algorithm Examples

To get the best results from the genetic algorithm, you usually need to experiment with different options. Selecting the best options for a problem involves trial and error. This section describes some ways you can change options to improve results. For a complete description of the available options, see “Importing and Exporting Options and Problems” on page 6-16.

This section covers the following topics:

- “Population Diversity” on page 7-34
- “Fitness Scaling” on page 7-39
- “Selection” on page 7-42
- “Reproduction Options” on page 7-43
- “Mutation and Crossover” on page 7-43
- “Setting the Amount of Mutation” on page 7-44
- “Setting the Crossover Fraction” on page 7-46
- “Comparing Results for Varying Crossover Fractions” on page 7-50
- “Global Versus Local Minima” on page 7-52
- “Using a Hybrid Function” on page 7-56
- “Setting the Maximum Number of Generations” on page 7-59
- “Vectorizing the Fitness Function” on page 7-60
- “Constrained Minimization Using ga” on page 7-61

### Population Diversity

One of the most important factors that determines the performance of the genetic algorithm performs is the *diversity* of the population. If the average distance between individuals is large, the diversity is high; if the average distance is small, the diversity is low. Getting the right amount of diversity is a matter of trial and error. If the diversity is too high or too low, the genetic algorithm might not perform well.

This section explains how to control diversity by setting the **Initial range** of the population. “Setting the Amount of Mutation” on page 7-44 describes how the amount of mutation affects diversity.

This section also explains how to set the population size.

### **Example – Setting the Initial Range**

By default, the Genetic Algorithm Tool creates a random initial population using the creation function. You can specify the range of the vectors in the initial population in the **Initial range** field in **Population** options.

---

**Note** The initial range only restricts the range of the points in the initial population by specifying the lower and upper bounds. Subsequent generations can contain points whose entries do not lie in the initial range. Additionally, the upper and lower bounds can be adjusted by setting the **Bounds** fields in the **Constraints** panel.

---

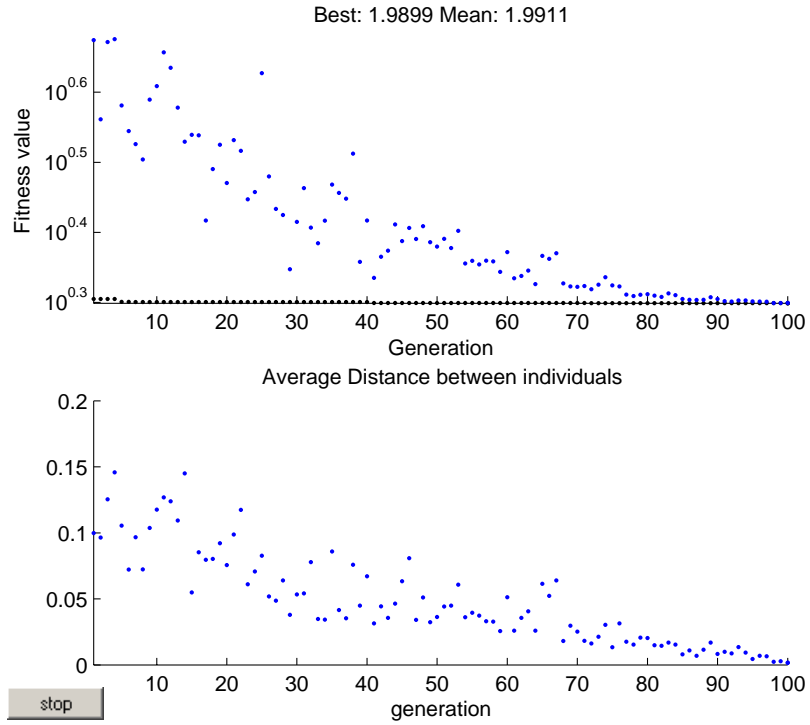
If you know approximately where the solution to a problem lies, you should specify the initial range so that it contains your guess for the solution. However, the genetic algorithm can find the solution even if it does not lie in the initial range, provided that the populations have enough diversity.

The following example shows how the initial range affects the performance of the genetic algorithm. The example uses Rastrigin’s function, described in “Example: Rastrigin’s Function” on page 4-7. The minimum value of the function is 0, which occurs at the origin.

To run the example, make the following settings in the Genetic Algorithm Tool:


- Set **Fitness function** to @rastriginsfcn.
- Set **Number of variables** to 2.
- Select **Best fitness** in the **Plots** pane.
- Select **Distance** in the **Plots** pane.
- Set **Initial range** to [1; 1.1].

Then click **Start**. The genetic algorithm returns the best fitness function value of approximately 2 and displays the plots in the following figure.




The upper plot, which displays the best fitness at each generation, shows little progress in lowering the fitness value. The lower plot shows the average distance between individuals at each generation, which is a good measure of the diversity of a population. For this setting of initial range, there is too little diversity for the algorithm to make progress.

Next, try setting **Initial range** to [1; 100] and running the algorithm. The genetic algorithm returns the best fitness value of approximately 3.9 and displays the following plots.



This time, the genetic algorithm makes progress, but because the average distance between individuals is so large, the best individuals are far from the optimal solution.

Finally, set **Initial range** to [1; 2] and run the genetic algorithm. This returns the best fitness value of approximately .012 and displays the following plots.



The diversity in this case is better suited to the problem, so the genetic algorithm returns a much better result than in the previous two cases.

### Setting the Population Size

The **Population size** field in **Population** options determines the size of the population at each generation. Increasing the population size enables the genetic algorithm to search more points and thereby obtain a better result. However, the larger the population size, the longer the genetic algorithm takes to compute each generation.

---

**Note** You should set **Population size** to be at least the value of **Number of variables**, so that the individuals in each population span the space being searched.

---

You can experiment with different settings for **Population size** that return good results without taking a prohibitive amount of time to run.

## Fitness Scaling

Fitness scaling converts the raw fitness scores that are returned by the fitness function to values in a range that is suitable for the selection function. The selection function uses the scaled fitness values to select the parents of the next generation. The selection function assigns a higher probability of selection to individuals with higher scaled values.


The range of the scaled values affects the performance of the genetic algorithm. If the scaled values vary too widely, the individuals with the highest scaled values reproduce too rapidly, taking over the population gene pool too quickly, and preventing the genetic algorithm from searching other areas of the solution space. On the other hand, if the scaled values vary only a little, all individuals have approximately the same chance of reproduction and the search will progress very slowly.

The default fitness scaling option, Rank, scales the raw scores based on the rank of each individual instead of its score. The rank of an individual is its position in the sorted scores: the rank of the most fit individual is 1, the next most fit is 2, and so on. The rank scaling function assigns scaled values so that


- The scaled value of an individual with rank  $n$  is proportional to  $1/(\sqrt{n})$ .
- The sum of the scaled values over the entire population equals the number of parents needed to create the next generation.

Rank fitness scaling removes the effect of the spread of the raw scores.

The following plot shows the raw scores of a typical population of 20 individuals, sorted in increasing order.



The following plot shows the scaled values of the raw scores using rank scaling.






Because the algorithm minimizes the fitness function, lower raw scores have higher scaled values. Also, because rank scaling assigns values that depend only on an individual's rank, the scaled values shown would be the same for any population of size 20 and number of parents equal to 32.


### Comparing Rank and Top Scaling

To see the effect of scaling, you can compare the results of the genetic algorithm using rank scaling with one of the other scaling options, such as Top. By default, top scaling assigns 40 percent of the fittest individuals to the same scaled value and assigns the rest of the individuals to value 0. Using the default selection function, only 40 percent of the fittest individuals can be selected as parents.

The following figure compares the scaled values of a population of size 20 with number of parents equal to 32 using rank and top scaling.



Because top scaling restricts parents to the fittest individuals, it creates less diverse populations than rank scaling. The following plot compares the variances of distances between individuals at each generation using rank and top scaling.



## Selection

The selection function chooses parents for the next generation based on their scaled values from the fitness scaling function. An individual can be selected more than once as a parent, in which case it contributes its genes to more than one child. The default selection option, `Stochastic uniform`, lays out a line in which each parent corresponds to a section of the line of length proportional to its scaled value. The algorithm moves along the line in steps of equal size. At each step, the algorithm allocates a parent from the section it lands on.

A more deterministic selection option is `Remainder`, which performs two steps:

- In the first step, the function selects parents deterministically according to the integer part of the scaled value for each individual. For example,

if an individual's scaled value is 2.3, the function selects that individual twice as a parent.

- In the second step, the selection function selects additional parents using the fractional parts of the scaled values, as in stochastic uniform selection. The function lays out a line in sections, whose lengths are proportional to the fractional part of the scaled value of the individuals, and moves along the line in equal steps to select the parents.

Note that if the fractional parts of the scaled values all equal 0, as can occur using Top scaling, the selection is entirely deterministic.

## Reproduction Options

Reproduction options control how the genetic algorithm creates the next generation. The options are

- **Elite count** — The number of individuals with the best fitness values in the current generation that are guaranteed to survive to the next generation. These individuals are called *elite children*. The default value of **Elite count** is 2.

When **Elite count** is at least 1, the best fitness value can only decrease from one generation to the next. This is what you want to happen, since the genetic algorithm minimizes the fitness function. Setting **Elite count** to a high value causes the fittest individuals to dominate the population, which can make the search less effective.

- **Crossover fraction** — The fraction of individuals in the next generation, other than elite children, that are created by crossover. “Setting the Crossover Fraction” on page 7-46 describes how the value of **Crossover fraction** affects the performance of the genetic algorithm.

## Mutation and Crossover

The genetic algorithm uses the individuals in the current generation to create the children that make up the next generation. Besides elite children, which correspond to the individuals in the current generation with the best fitness values, the algorithm creates

- Crossover children by selecting vector entries, or genes, from a pair of individuals in the current generation and combines them to form a child

- Mutation children by applying random changes to a single individual in the current generation to create a child

Both processes are essential to the genetic algorithm. Crossover enables the algorithm to extract the best genes from different individuals and recombine them into potentially superior children. Mutation adds to the diversity of a population and thereby increases the likelihood that the algorithm will generate individuals with better fitness values.

See “Creating the Next Generation” on page 4-20 for an example of how the genetic algorithm applies mutation and crossover.

You can specify how many of each type of children the algorithm creates as follows:

- **Elite count**, in **Reproduction** options, specifies the number of elite children.
- **Crossover fraction**, in **Reproduction** options, specifies the fraction of the population, other than elite children, that are crossover children.

For example, if the **Population size** is 20, the **Elite count** is 2, and the **Crossover fraction** is 0.8, the numbers of each type of children in the next generation are as follows:

- There are two elite children.
- There are 18 individuals other than elite children, so the algorithm rounds  $0.8 \times 18 = 14.4$  to 14 to get the number of crossover children.
- The remaining four individuals, other than elite children, are mutation children.


### Setting the Amount of Mutation

The genetic algorithm applies mutations using the option that you specify on the **Mutation function** pane. The default mutation option, **Gaussian**, adds a random number, or *mutation*, chosen from a Gaussian distribution, to each entry of the parent vector. Typically, the amount of mutation, which is proportional to the standard deviation of the distribution, decreases at each new generation. You can control the average amount of mutation that

the algorithm applies to a parent in each generation through the **Scale** and **Shrink** options:

- **Scale** controls the standard deviation of the mutation at the first generation, which is **Scale** multiplied by the range of the initial population, which you specify by the **Initial range** option.
- **Shrink** controls the rate at which the average amount of mutation decreases. The standard deviation decreases linearly so that its final value equals  $1 - \text{Shrink}$  times its initial value at the first generation. For example, if **Shrink** has the default value of 1, then the amount of mutation decreases to 0 at the final step.


You can see the effect of mutation by selecting the plot options **Distance** and **Range**, and then running the genetic algorithm on a problem such as the one described in “Example: Rastrigin’s Function” on page 4-7. The following figure shows the plot.



The upper plot displays the average distance between points in each generation. As the amount of mutation decreases, so does the average distance between individuals, which is approximately 0 at the final generation. The

lower plot displays a vertical line at each generation, showing the range from the smallest to the largest fitness value, as well as mean fitness value. As the amount of mutation decreases, so does the range. These plots show that reducing the amount of mutation decreases the diversity of subsequent generations.

For comparison, the following figure shows the plots for **Distance** and **Range** when you set **Shrink** to 0.5.



With **Shrink** set to 0.5, the average amount of mutation decreases by a factor of 1/2 by the final generation. As a result, the average distance between individuals decreases by approximately the same factor.

## Setting the Crossover Fraction

The **Crossover fraction** field, in the **Reproduction** options, specifies the fraction of each population, other than elite children, that are made up of crossover children. A crossover fraction of 1 means that all children other than

elite individuals are crossover children, while a crossover fraction of 0 means that all children are mutation children. The following example show that neither of these extremes is an effective strategy for optimizing a function.

The example uses the fitness function whose value at a point is the sum of the absolute values of the coordinates at the points. That is,

$$f(x_1, x_2, \dots, x_n) = |x_1| + |x_2| + \dots + |x_n|$$


You can define this function as an anonymous function by setting **Fitness function** to

```
@(x) sum(abs(x))
```

To run the example,

- Set **Fitness function** to `@(x) sum(abs(x))`.
- Set **Number of variables** to 10.
- Set **Initial range** to `[-1; 1]`.
- Select **Best fitness** and **Distance** in the **Plots** pane.


Run the example with the default value of 0.8 for **Crossover fraction**. This returns the best fitness value of approximately 0.2 and displays the following plots.





## Crossover Without Mutation


To see how the genetic algorithm performs when there is no mutation, set **Crossover fraction** to 1.0 and click **Start**. This returns the best fitness value of approximately 1.3 and displays the following plots.



In this case, the algorithm selects genes from the individuals in the initial population and recombines them. The algorithm cannot create any new genes because there is no mutation. The algorithm generates the best individual that it can using these genes at generation number 8, where the best fitness plot becomes level. After this, it creates new copies of the best individual, which are then selected for the next generation. By generation number 17, all individuals in the population are the same, namely, the best individual. When this occurs, the average distance between individuals is 0. Since the algorithm cannot improve the best fitness value after generation 8, it stalls after 50 more generations, because **Stall generations** is set to 50.

### Mutation Without Crossover

To see how the genetic algorithm performs when there is no crossover, set **Crossover fraction** to 0 and click **Start**. This returns the best fitness value of approximately 3.5 and displays the following plots.



In this case, the random changes that the algorithm applies never improve the fitness value of the best individual at the first generation. While it improves the individual genes of other individuals, as you can see in the upper plot by the decrease in the mean value of the fitness function, these improved genes are never combined with the genes of the best individual because there is no crossover. As a result, the best fitness plot is level and the algorithm stalls at generation number 50.

### Comparing Results for Varying Crossover Fractions


The demo `deterministicstudy.m`, which is included in the toolbox, compares the results of applying the genetic algorithm to Rastrigin's function with

**Crossover fraction** set to 0, .2, .4, .6, .8, and 1. The demo runs for 10 generations. At each generation, the demo plots the means and standard deviations of the best fitness values in all the preceding generations, for each value of the **Crossover fraction**.

To run the demo, enter

```
deterministicstudy
```

at the MATLAB prompt. When the demo is finished, the plots appear as in the following figure.



The lower plot shows the means and standard deviations of the best fitness values over 10 generations, for each of the values of the crossover fraction. The upper plot shows a color-coded display of the best fitness values in each generation.

For this fitness function, setting **Crossover fraction** to 0.8 yields the best result. However, for another fitness function, a different setting for **Crossover fraction** might yield the best result.


## Global Versus Local Minima

Sometimes the goal of an optimization is to find the global minimum or maximum of a function—a point where the function value is smaller or larger at any other point in the search space. However, optimization algorithms sometimes return a local minimum—a point where the function value is smaller than at nearby points, but possibly greater than at a distant point in the search space. The genetic algorithm can sometimes overcome this deficiency with the right settings.

As an example, consider the following function

$$f(x) = \begin{cases} -\exp\left(-\left(\frac{x}{20}\right)^2\right) & \text{for } x \leq 20 \\ -\exp(-1) + (x - 20)(x - 22) & \text{for } x > 20 \end{cases}$$

The following figure shows a plot of the function.



The function has two local minima, one at  $x = 0$ , where the function value is  $-1$ , and the other at  $x = 21$ , where the function value is  $-1 - 1/e$ . Since the latter value is smaller, the global minimum occurs at  $x = 21$ .

## Running the Genetic Algorithm on the Example

To run the genetic algorithm on this example,

**1** Copy and paste the following code into a new M-file in the MATLAB Editor.

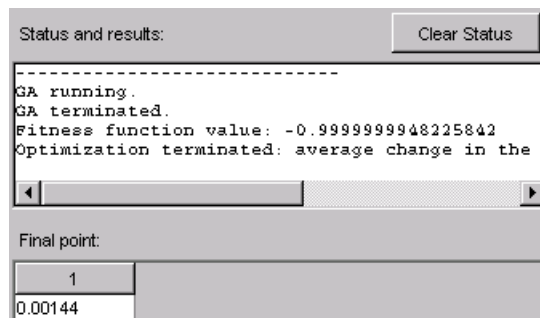
```
function y = two_min(x)
if x<=20
    y = -exp(-(x/20).^2);
else
    y = -exp(-1)+(x-20)*(x-22);
end
```

**2** Save the file as `two_min.m` in a directory on the MATLAB path.


**3** In the Genetic Algorithm Tool,

- Set **Fitness function** to `@two_min`.
- Set **Number of variables** to 1.
- Click **Start**.

The genetic algorithm returns a point very close to the local minimum at  $x = 0$ .




The following custom plot shows why the algorithm finds the local minimum rather than the global minimum. The plot shows the range of individuals in each generation and the best individual.



Note that all individuals are between -2 and 2.5. While this range is larger than the default **Initial range** of [0; 1], due to mutation, it is not large enough to explore points near the global minimum at  $x = 21$ .


One way to make the genetic algorithm explore a wider range of points—that is, to increase the diversity of the populations—is to increase the **Initial range**. The **Initial range** does not have to include the point  $x = 21$ , but it must be large enough so that the algorithm generates individuals near  $x = 21$ . Set **Initial range** to  $[0; 15]$  as shown in the following figure.



The screenshot shows a dialog box titled "Population" with the following settings:

- Population type: Double Vector
- Population size: 20
- Creation function: Uniform
- Initial population: 0
- Initial scores: 0
- Initial range: [0 ; 15]

Then click **Start**. The genetic algorithm returns a point very close to 21.




The screenshot shows a dialog box titled "Status and results" with a "Clear Status" button. The text area contains the following output:

```
-----  
GA running.  
GA terminated.  
Fitness function value: -1.3664110605692392  
Optimization terminated: average change in the :  
-----
```

Below the text area is a horizontal scrollbar. Underneath, the "Final point:" section displays a table with the following data:

| 1        |
|----------|
| 21.03832 |

This time, the custom plot shows a much wider range of individuals. By the second generation there are individuals greater than 21, and by generation 12, the algorithm finds a best individual that is approximately equal to 21.



## Using a Hybrid Function

A hybrid function is an optimization function that runs after the genetic algorithm terminates in order to improve the value of the fitness function. The hybrid function uses the final point from the genetic algorithm as its initial point. You can specify a hybrid function in **Hybrid function** options.


This example uses the function `fminunc`, an unconstrained minimization function in the Optimization Toolbox. The example first runs the genetic algorithm to find a point close to the optimal point and then uses that point as the initial point for `fminunc`.

The example finds the minimum of Rosenbrock's function, which is defined by

$$f(x_1, x_2) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2$$



The following figure shows a plot of Rosenbrock's function.



The toolbox provides an M-file, `dejong2fcn.m`, that computes the function. To see a demo of this example, enter


```
hybriddemo
```

at the MATLAB prompt.


To explore the example, first enter `gatool` to open the Genetic Algorithm Tool and enter the following settings:

- Set **Fitness function** to `@dejong2fcn`.
- Set **Number of variables** to 2.
- Set **Population size** to 10.

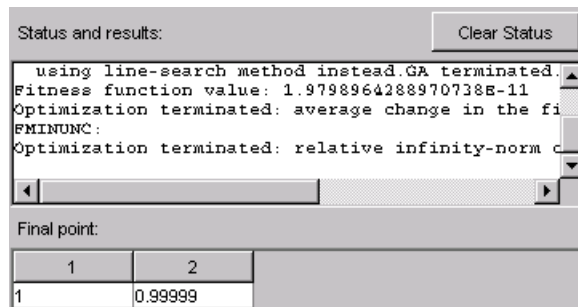
Before adding a hybrid function, try running the genetic algorithm by itself, by clicking **Start**. The genetic algorithm displays the following results in the **Status and results** pane.



The final point is close to the true minimum at (1, 1). You can improve this result by setting **Hybrid function** to **fminunc** in the **Hybrid function** options.



The function **fminunc** uses the final point of one run, or step, of the genetic algorithm as the initial point of the next step to return a more accurate result when the genetic algorithm terminates, as shown in the **Status and results** pane.




## Setting the Maximum Number of Generations

The **Generations** option in **Stopping criteria** determines the maximum number of generations the genetic algorithm runs for—see “Stopping Conditions for the Algorithm” on page 4-22. Increasing the **Generations** option often improves the final result.

As an example, change the settings in the Genetic Algorithm Tool as follows:

- Set **Fitness function** to @rastriginsfcn.
- Set **Number of variables** to 10.
- Select **Best fitness** in the **Plots** pane.
- Set **Generations** to Inf.
- Set **Stall generations** to Inf.
- Set **Stall time** to Inf.

Then run the genetic algorithm for approximately 300 generations and click **Stop**. The following figure shows the resulting best fitness plot after 300 generations.



Note that the algorithm *stalls* at approximately generation number 170—that is, there is no immediate improvement in the fitness function after generation 170. If you restore **Stall generations** to its default value of 50, the algorithm would terminate at approximately generation number 230. If the genetic algorithm stalls repeatedly with the current setting for **Generations**, you can try increasing both the **Generations** and **Stall generations** options to improve your results. However, changing other options might be more effective.

---

**Note** When **Mutation function** is set to Gaussian, increasing the value of **Generations** might actually worsen the final result. This can occur because the Gaussian mutation function decreases the average amount of mutation in each generation by a factor that depends on the value specified in **Generations**. Consequently, the setting for **Generations** affects the behavior of the algorithm.

---

## Vectorizing the Fitness Function

The genetic algorithm usually runs faster if you *vectorize* the fitness function. This means that the genetic algorithm only calls the fitness function once, but expects the fitness function to compute the fitness for all individuals in the current population at once. To vectorize the fitness function,

- Write the M-file that computes the function so that it accepts a matrix with arbitrarily many rows, corresponding to the individuals in the population. For example, to vectorize the function

$$f(x_1, x_2) = x_1^2 - 2x_1x_2 + 6x_1 + x_2^2 - 6x_2$$

write the M-file using the following code:

```
z = x(:,1).^2 - 2*x(:,1).*x(:,2) + 6*x(:,1) + x(:,2).^2 - 6*x(:,2);
```

The colon in the first entry of  $x$  indicates all the rows of  $x$ , so that  $x(:, 1)$  is a vector. The  $\wedge$  and  $\cdot$  operators perform element-wise operations on the vectors.

- In the **Vectorize** pane, set the **Fitness function is vectorized** option to On.

---

**Note** The fitness function must accept an arbitrary number of rows to use the **Vectorize** option.

---

The following comparison, run at the command line, shows the improvement in speed with **Vectorize** set to `On`.

```
tic;ga(@rastriginsfcn,20);toc

elapsed_time =

    4.3660
options=gaoptimset('Vectorize','on');
tic;ga(@rastriginsfcn,20,[],[],[],[],[],[],[],options);toc

elapsed_time =

    0.5810
```

## Constrained Minimization Using `ga`

Suppose you want to minimize the simple fitness function of two variables  $x_1$  and  $x_2$ ,

$$\min_x f(x) = 100 (x_1^2 - x_2)^2 + (1 - x_1)^2$$

subject to the following nonlinear inequality constraints and bounds

$$\begin{aligned} x_1 \cdot x_2 + x_1 - x_2 + 1.5 &\leq 0 && \text{(nonlinear constraint)} \\ 10 - x_1 \cdot x_2 &\leq 0 && \text{(nonlinear constraint)} \\ 0 &\leq x_1 \leq 1 && \text{(bound)} \\ 0 &\leq x_2 \leq 13 && \text{(bound)} \end{aligned}$$

Begin by creating the fitness and constraint functions. First, create an M-file named `simple_fitness.m` as follows:

```
function y = simple_fitness(x)
```

$$y = 100 * (x(1)^2 - x(2)) ^2 + (1 - x(1))^2;$$

The genetic algorithm function, `ga`, assumes the fitness function will take one input `x`, where `x` has as many elements as the number of variables in the problem. The fitness function computes the value of the function and returns that scalar value in its one return argument, `y`.

Then create an M-file, `simple_constraint.m`, containing the constraints

```
function [c, ceq] = simple_constraint(x)
c = [1.5 + x(1)*x(2) + x(1) - x(2);
-x(1)*x(2) + 10];
ceq = [];
```

The `ga` function assumes the constraint function will take one input `x`, where `x` has as many elements as the number of variables in the problem. The constraint function computes the values of all the inequality and equality constraints and returns two vectors, `c` and `ceq`, respectively.

To minimize the fitness function, you need to pass a function handle to the fitness function as the first argument to the `ga` function, as well as specifying the number of variables as the second argument. Lower and upper bounds are provided as `LB` and `UB` respectively. In addition, you also need to pass a function handle to the nonlinear constraint function.

```
ObjectiveFunction = @simple_fitness;
nvars = 2; % Number of variables
LB = [0 0]; % Lower bound
UB = [1 13]; % Upper bound
ConstraintFunction = @simple_constraint;
[x,fval] = ga(ObjectiveFunction,nvars,[],[],[],[],LB,UB,ConstraintFunction)
```

```
Warning: 'mutationgaussian' mutation function
is for unconstrained minimization only;
using @mutationadaptfeasible mutation function.
Set @mutationadaptfeasible as MutationFcn options
using GAOPTIMSET.
```

```
Optimization terminated: current tolerance on f(x) 1e-007
is less than options.TolFun and constraint violation is
```

```
less than options.TolCon.
```

```
x =
```

```
    0.8122    12.3122
```

```
fval =
```

```
1.3578e+004
```

---

**Note** For the constrained minimization problem, the `ga` function changed the mutation function to `@mutationadaptfeasible`. The default mutation function, `@mutationgaussian`, is only appropriate for unconstrained minimization problems.

---

The genetic algorithm solver handles linear constraints and bounds differently from nonlinear constraints. All the linear constraints and bounds are satisfied throughout the optimization. However, `ga` may not satisfy all the nonlinear constraints at every generation. If `ga` converges to a solution, the nonlinear constraints will be satisfied at that solution.

`ga` uses the mutation and crossover functions to produce new individuals at every generation. `ga` satisfies linear and bound constraints by using mutation and crossover functions that only generate feasible points. For example, in the previous call to `ga`, the default mutation function `mutationgaussian` will not satisfy the linear constraints and so the `mutationadaptfeasible` is used instead. If you provide a custom mutation function, this custom function must only generate points that are feasible with respect to the linear and bound constraints. All the crossover functions in the toolbox generate points that satisfy the linear constraints and bounds except the `crossoverheuristic` function.

Specify `mutationadaptfeasible` as the mutation function for the minimization problem by using the `gaoptimset` function.

```
options = gaoptimset('MutationFcn',@mutationadaptfeasible);
```

Next run the ga solver.

```
[x,fval] = ga(ObjectiveFunction,nvars,[],[],[],[],LB,UB,ConstraintFunction,options)
```

```
Optimization terminated: current tolerance on f(x) 1e-007 is less than options.TolFun  
and constraint violation is less than options.TolCon.
```

```
x =
```

```
0.8122 12.3122
```

```
fval =
```

```
1.3578e+004
```

Now, use the `gaoptimset` function to create an options structure that will select two plot functions. The first plot function is `gaplotbestf`, which plots the best and mean score of the population at every generation. The second plot function is `gaplotmaxconstr`, which plots the maximum constraint violation of nonlinear constraints at every generation. You can also visualize the progress of the algorithm by displaying information to the command window using the `'Display'` option.

```
options = gaoptimset(options,'PlotFcns',{@gaplotbestf,@gaplotmaxconstr},'Display','iter');
```



Rerun the ga solver.

```
[x,fval] = ga(ObjectiveFunction,nvars,[],[],[],[],...
             LB,UB,ConstraintFunction,options)
```

| Generation | f-count | Best<br>f(x) | max<br>constraint | Stall<br>Generations |
|------------|---------|--------------|-------------------|----------------------|
| 1          | 849     | 14915.8      | 0                 | 0                    |
| 2          | 1567    | 13578.3      | 0                 | 0                    |
| 3          | 2334    | 13578.3      | 0                 | 1                    |
| 4          | 3043    | 13578.3      | 0                 | 2                    |
| 5          | 3752    | 13578.3      | 0                 | 3                    |

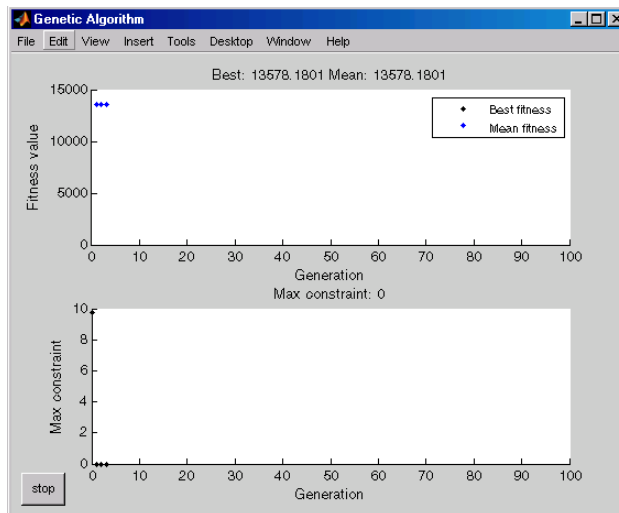
Optimization terminated: current tolerance on f(x) 1e-009  
is less than options.TolFun and constraint violation is  
less than options.TolCon.

x =

```
0.8122 12.3123
```

fval =

```
1.3578e+004
```



You can provide a start point for the minimization to the ga function by specifying the `InitialPopulation` option. The ga function will use the first individual from `InitialPopulation` as a start point for a constrained minimization.

```
X0 = [0.5 0.5]; % Start point (row vector)
options = gaoptimset(options,'InitialPopulation',X0);
```

Now, rerun the ga solver.

```
[x,fval] = ga(ObjectiveFunction,nvars,[],[],[],[],...
              LB,UB,ConstraintFunction,options)
```

| Generation | f-count | Best<br>f(x) | max<br>constraint | Stall<br>Generations |
|------------|---------|--------------|-------------------|----------------------|
| 1          | 965     | 13579.6      | 0                 | 0                    |
| 2          | 1728    | 13578.2      | 1.776e-015        | 0                    |
| 3          | 2422    | 13578.2      | 0                 | 0                    |

Optimization terminated: current tolerance on f(x) 1e-007  
is less than options.TolFun and constraint violation is  
less than options.TolCon.

```
x =  
  
    0.8122    12.3122  
  
fval =  
  
    1.3578e+004
```

## Parameterizing Functions Called by ga

Sometimes you might want to write functions that are called by `ga` that have additional parameters to the independent variable. For example, suppose you want to minimize the following function:

$$f(x) = (a - bx_1^2 + x_1^{4/3})x_1^2 + x_1x_2 + (-c + cx_2^2)x_2^2$$

for different values of  $a$ ,  $b$ , and  $c$ . Because `ga` accepts a fitness function that depends only on  $x$ , you must provide the additional parameters  $a$ ,  $b$ , and  $c$  to the function before calling `ga`. The following sections describe two ways to do this:

- “Parameterizing Functions Using Anonymous Functions with `ga`” on page 7-67
- “Parameterizing a Function Using a Nested Function with `ga`” on page 7-68

The examples in these sections show how to parameterize the objective function, but you can use the same methods to parameterize any user-defined functions called by `ga`, `patternsearch`, `simulannealbnd`, or `thresacceptbnd`.

### Parameterizing Functions Using Anonymous Functions with `ga`

To parameterize your function, first write an M-file containing the following code:

```
function y = parameterfun(x,a,b,c)
y = (a - b*x(1)^2 + x(1)^4/3)*x(1)^2 + x(1)*x(2) + ...
    (-c + c*x(2)^2)*x(2)^2;
```

Save the M- file as `parameterfun.m` in a directory on the MATLAB path.

Now, suppose you want to minimize the function for the parameter values  $a = 4$ ,  $b = 2.1$ , and  $c = 4$ . To do so, define a function handle to an anonymous function by entering the following commands at the MATLAB prompt:

```
a = 4; b = 2.1; c = 4;    % Define parameter values
fitfun = @(x) parameterfun(x,a,b,c);
NVAR = 3;
```

If you are using the Genetic AlgorithmTool,

- Set **Fitness function** to fitfun.
- Set **Number of variables** to NVARS.

```
a = 3.6; b = 2.4; c = 5;    % Define parameter values
fitfun = @(x) parameterfun(x,a,b,c);
```

## Parameterizing a Function Using a Nested Function with ga

As an alternative to parameterizing the objective function as an anonymous function, you can write a single M-file that

- Accepts  $a$ ,  $b$ ,  $c$  and  $NVAR$  as inputs.
- Contains the fitness function as a nested function.
- Calls `ga`.

The following shows the code for the M-file.

```
function [x fval] = runga(a,b,c,NVAR)
[x, fval] = ga(@nestedfun,NVAR);
% Nested function that computes the fitness function
function y = nestedfun(x)
    y = (a - b*x(1)^2 + x(1)^4/3)*x(1)^2 + x(1)*x(2) + ...
        (-c + c*x(2)^2)*x(2)^2;
end
end
```

Note that the fitness function is computed in the nested function `nestedfun`, which has access to the variables `a`, `b`, and `c`. To run the optimization, enter

```
[x fval] = runga(a,b,c,NVARS)
```

This returns

```
Optimization terminated: average change in the fitness  
value less than options.TolFun.
```

```
x =
```

```
-0.1302    0.7170    0.2272
```

```
fval =
```

```
-1.0254
```



# Using the Simulated Annealing and Threshold Acceptance Algorithms

---

Using the Simulated Annealing and Threshold Acceptance Algorithms from the Command Line (p. 8-2)

Explains how to use the simulated annealing and threshold acceptance algorithms at the command line.

Simulated Annealing and Threshold Acceptance Examples (p. 8-7)

Explains how to set options for the simulated annealing and threshold acceptance algorithms.

## Using the Simulated Annealing and Threshold Acceptance Algorithms from the Command Line

You can run the simulated annealing function `simulannealbnd` and threshold acceptance function `threshacceptbnd` from the command line. This section explains how to do so and covers the following topics:

- “Running `simulannealbnd` and `threshacceptbnd` with the Default Options” on page 8-2
- “Setting Options for `simulannealbnd` and `threshacceptbnd` at the Command Line” on page 8-3
- “Reproducing Your Results” on page 8-5

### Running `simulannealbnd` and `threshacceptbnd` with the Default Options

To run either the simulated annealing or threshold acceptance algorithms with the default options, call the corresponding function with the syntax

```
[x, fval] = simulannealbnd(@objfun, x0)
[x, fval] = threshacceptbnd(@objfun, x0)
```

The input arguments to `simulannealbnd` and `threshacceptbnd` are

- `@objfun` — A function handle to the M-file that computes the objective function. Chapter 2, “Writing M-Files for Functions You Want to Optimize” explains how to write this M-file.
- `x0` — The initial guess of the optimal argument to the objective function.

The output arguments are

- `x` — The final point.
- `fval` — The value of the objective function at `x`.

For a description of additional input and output arguments, see the reference pages for `simulannealbnd` and `threshacceptbnd`.



You can run the example described in “Example: Minimizing De Jong’s Fifth Function” on page 5-4 from the command line with the simulated annealing algorithm by entering

```
[x, fval] = simulannealbnd(@dejong5fcn, [0 0])
```

This returns

```
x =  
   -31.9564   -15.9755  
fval =  
    5.9288
```

### Additional Output Arguments

To get more information about the performance of the algorithm, you can call `simulannealbnd` or `threshacceptbnd` with the syntax

```
[x, fval, exitflag, output] = simulannealbnd(@objfun, x0)
```

Besides `x` and `fval`, this function returns the following additional output arguments:

- `exitflag` — Flag indicating the reason the algorithm terminated
- `output` — Structure containing information about the performance of the algorithm

See the `simulannealbnd` and `threshacceptbnd` reference pages for more information about these arguments.

### Setting Options for `simulannealbnd` and `threshacceptbnd` at the Command Line

You can specify options by passing an options structure as an input argument to either `simulannealbnd` and `threshacceptbnd` using the syntax

```
[x, fval] = simulannealbnd(@objfun, x0, [], [], options)
```

This syntax does not specify any lower or upper bound constraints.

You create the options structure using the function `saoptimset`:

```
options = saoptimset('simulannealbnd')
```

This returns the structure options with the default values for its fields:

```
options =  
    AnnealingFcn: @annealingfast  
    TemperatureFcn: @temperatureexp  
    AcceptanceFcn: @acceptancesa  
        TolFun: 1.0000e-006  
    StallIterLimit: '500*numberofvariables'  
    MaxFunEvals: '3000*numberofvariables'  
    TimeLimit: Inf  
    MaxIter: Inf  
    ObjectiveLimit: -Inf  
    Display: 'final'  
    DisplayInterval: 10  
    HybridFcn: []  
    HybridInterval: 'end'  
    PlotFcns: []  
    PlotInterval: 1  
    OutputFcns: []  
    InitialTemperature: 100  
    ReannealInterval: 100  
    DataType: 'double'
```

These are the default values for simulannealbnd. To see the default values for threshacceptbnd, run

```
options = saoptimset('threshacceptbnd')
```

The value of each option is stored in a field of the options structure, such as options.ReannealInterval. You can display any of these values by entering options followed by the name of the field. For example, to display the interval for reannealing used for the simulated annealing algorithm, enter

```
options.ReannealInterval  
ans =  
    100
```

To create an options structure with a field value that is different from the default—for example, to set `ReannealInterval` to 300 instead of its default value 100—enter

```
options = saoptimset('ReannealInterval', 300)
```

This creates the options structure with all values set to their defaults, except for `ReannealInterval`, which is set to 300.

If you now enter

```
simulannealbnd(@objfun, x0, [], [], options)
```

`simulannealbnd` runs the simulated annealing algorithm with a reannealing interval of 300.

If you subsequently decide to change another field in the options structure, such as setting `PlotFcns` to `@splotbestf`, which plots the best objective function value at each iteration, call `saoptimset` with the syntax

```
options = saoptimset(options, 'PlotFcns', @splotbestf)
```

This preserves the current values of all fields of options except for `PlotFcns`, which is changed to `@splotbestf`. Note that if you omit the input argument `options`, `saoptimset` resets `ReannealInterval` to its default value 100.

You can also set both `ReannealInterval` and `PlotFcns` with the single command

```
options = saoptimset('ReannealInterval', 300, ...  
                    'PlotFcns', @splotbestf)
```

## Reproducing Your Results

Because the simulated annealing and threshold acceptance algorithms are stochastic—that is, they each makes random choices—you get slightly different results each time you run them. The algorithms use the MATLAB uniform and normal random number generators, `rand` and `randn`, when generating subsequent points and also when determining whether or not to accept new points. Each time the algorithms call `rand` and `randn`, their states are changed so that the next time they are called, they return different random numbers.

If you need to reproduce your results exactly, call `simulannealbnd` or `threshacceptbnd` with an output argument that contains the current states of `rand` and `randn` and then reset the states to these values before running the function again. For example, to reproduce the output of `simulanneal` applied to De Jong's fifth function, call `simulannealbnd` with the syntax

```
[x,fval,exitflag,output] = simulannealbnd(@dejong5fcn,[0 0]);
```

Suppose the results are

```
x =  
-32.0401 -16.1223  
fval =  
5.9288
```

The states of `rand` and `randn` are stored in two fields of `output`.

```
output =  
  iterations: 2041  
  funccount: 2058  
  message: [1x80 char]  
  randstate: [625x1 uint32]  
  randnstate: [2x1 double]  
  problemtype: 'unconstrained'  
  temperature: [2x1 double]  
  totaltime: 1.8226
```

Reset the states by entering

```
rand('twister', output.randstate);  
randn('state', output.randnstate);
```

If you now run `simulannealbnd` a second time, you get the same results.

---

**Note** If you do not need to reproduce your results, it is better not to set the states of `rand` and `randn`, so that you get the benefit of the randomness in these algorithms.

---

## **Simulated Annealing and Threshold Acceptance Examples**

If you are viewing this documentation in the Help browser, see the demo [Minimization Using Simulated Annealing And Threshold Acceptance Algorithms](#).



# Options Reference

---

Pattern Search Options (p. 9-2)

Describes the options for pattern search.

Genetic Algorithm Options (p. 9-23)

Describes the options for the genetic algorithm.

Simulated Annealing and Threshold  
Acceptance Algorithm Options  
(p. 9-45)

Describes the options for the simulated annealing and threshold acceptance algorithms.

## Pattern Search Options

This section describes the options for pattern search. There are two ways to specify the options, depending on whether you are using the Pattern Search Tool or calling the function `patternsearch` at the command line:

- If you are using the Pattern Search Tool (`psearchtool`), you specify the options by selecting an option from a drop-down list or by entering the value of the option in the text field, as described in “Setting Options in the Pattern Search Tool” on page 6-15.
- If you are calling `patternsearch` from the command line, you specify the options by creating an options structure using the function `psoptimset`, as follows:

```
options = psoptimset('Param1', value1, 'Param2', value2, ...);
```

See “Setting Options for `patternsearch` at the Command Line” on page 6-22 for examples.

In this section, each option is listed in two ways:

- By its label, as it appears in the Pattern Search Tool
- By its field name in the `options` structure

For example:

- **Poll method** refers to the label of the option in the Pattern Search Tool.
- `PollMethod` refers to the corresponding field of the `options` structure.

The options are divided into the following categories:

- “Plot Options” on page 9-3
- “Poll Options” on page 9-5
- “Search Options” on page 9-8
- “Mesh Options” on page 9-12
- “Algorithm Settings” on page 9-13



- “Cache Options” on page 9-14
- “Stopping Criteria” on page 9-14
- “Output Function Options” on page 9-16
- “Display to Command Window Options” on page 9-18
- “Vectorize Option” on page 9-18
- “Options Table for Pattern Search Algorithms” on page 9-19

## Plot Options

Plot options enable you to plot data from the pattern search while it is running. When you select plot functions and run the pattern search, a plot window displays the plots on separate axes. You can stop the algorithm at any time by clicking the **Stop** button on the plot window.

**Plot interval** (`PlotInterval`) specifies the number of iterations between consecutive calls to the plot function.

You can select any of the following plots in the **Plots** pane.

- **Best function value** (`@psplotbestf`) plots the best objective function value.
- **Function count** (`@psplotfuncount`) plots the number of function evaluations.
- **Mesh size** (`@psplotmeshsize`) plots the mesh size.
- **Best point** (`@psplotbestx`) plots the current best point.
- **Max constraint** (`@psplotmaxconstr`) plots the maximum nonlinear constraint violation.
- **Custom** enables you to use your own plot function. To specify the plot function using the Pattern Search Tool,
  - Select **Custom function**.
  - Enter `@myfun` in the text box, where `myfun` is the name of your function.

“Structure of the Plot Functions” on page 9-4 describes the structure of a plot function.

To display a plot when calling `patternsearch` from the command line, set the `PlotFcns` field of options to be a function handle to the plot function. For example, to display the best function value, set options as follows

```
options = psoptimset('PlotFcns', @psplotbestf);
```

To display multiple plots, use the syntax

```
options = psoptimset('PlotFcns', {@plotfun1, @plotfun2, ...});
```

where `@plotfun1`, `@plotfun2`, and so on are function handles to the plot functions (listed in parentheses in the preceding list).

### Structure of the Plot Functions

The first line of a plot function has the form

```
function stop = plotfun(optimvalues, flag)
```

The input arguments to the function are

- `optimvalues` — Structure containing information about the current state of the solver. The structure contains the following fields:
  - `x` — Current point
  - `iteration` — Iteration number
  - `fval` — Objective function value
  - `meshsize` — Current mesh size
  - `funccount` — Number of function evaluations
  - `method` — Method used in last iteration
  - `TolFun` — Tolerance on function value in last iteration
  - `TolX` — Tolerance on `x` value in last iteration
  - `nonlineq` — Nonlinear inequality constraints, displayed only when a nonlinear constraint function is specified
  - `nonlineq` — Nonlinear equality constraints, displayed only when a nonlinear constraint function is specified

- `flag` — Current state in which the plot function is called. The possible values for `flag` are
  - `init` — Initialization state
  - `iter` — Iteration state
  - `interrupt` — Intermediate stage
  - `done` — Final state

“Parameterizing Functions Called by `patternsearch`” on page 6-53 explains how to provide additional parameters to the function.

The output argument `stop` provides a way to stop the algorithm at the current iteration. `stop` can have the following values:

- `false` — The algorithm continues to the next iteration.
- `true` — The algorithm terminates at the current iteration.

## Poll Options

Poll options control how the pattern search polls the mesh points at each iteration.

**Poll method** (`PollMethod`) specifies the pattern the algorithm uses to create the mesh. There are two patterns for each of the two classes of direct search algorithms: the generalized pattern search (GPS) algorithm and the mesh adaptive direct search (MADS) algorithm.

These patterns are the Positive basis  $2N$  and the Positive basis  $N+1$ :

- The default pattern, GPS Positive basis  $2N$ , consists of the following  $2N$  vectors, where  $N$  is the number of independent variables for the objective function.

```
[100...0]
[010...0]
...
[000...1]
[-1 00...0]
[0 -10...0]
...
[0 00...-1]
```

For example, if the optimization problem has three independent variables, the pattern consists of the following six vectors.

```
[1 0 0]
[0 1 0]
[0 0 1]
[-1 0 0]
[0 -1 0]
[0 0 -1]
```

- The pattern, MADS Positive basis  $2N$ , consists of  $2N$  randomly generated vectors, where  $N$  is the number of independent variables for the objective function. This is done by randomly generating  $N$  vectors which form a linearly independent set, then using this first set and the negative of this set gives  $2N$  vectors. As shown above, the GPS Positive basis  $2N$  pattern is formed using the positive and negative of the linearly independent identity, however, with the MADS Positive basis  $2N$ , the pattern is generated using a random permutation of an  $N$ -by- $N$  linearly independent lower triangular matrix that is regenerated at each iteration.

- The GPS Positive basis NP1 pattern consists of the following  $N + 1$  vectors.

$$\begin{bmatrix} 100\dots0 \\ 010\dots0 \\ \dots \\ 000\dots1 \\ -1-1-1\dots-1 \end{bmatrix}$$

For example, if the objective function has three independent variables, the pattern consists of the following four vectors.

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ -1 & -1 & -1 \end{bmatrix}$$

- The pattern, MADS Positive basis N+1, consists of N randomly generated vectors to form the positive basis, where N is the number of independent variables for the objective function. Then, one more random vector is generated, giving N+1 randomly generated vectors. Each iteration generates a new pattern when the MADS Positive basis N+1 is selected.

**Complete poll** (CompletePoll) specifies whether all the points in the current mesh must be polled at each iteration. **Complete Poll** can have the values On or Off.

- If you set **Complete poll** to On, the algorithm polls all the points in the mesh at each iteration and chooses the point with the smallest objective function value as the current point at the next iteration.
- If you set **Complete poll** to Off, the default value, the algorithm stops the poll as soon as it finds a point whose objective function value is less than that of the current point. The algorithm then sets that point as the current point at the next iteration.

**Polling order** (PollingOrder) specifies the order in which the algorithm searches the points in the current mesh. The options are

- Random — The polling order is random.
- Success — The first search direction at each iteration is the direction in which the algorithm found the best point at the previous iteration. After the first point, the algorithm polls the mesh points in the same order as **Consecutive**.
- Consecutive — The algorithm polls the mesh points in *consecutive* order, that is, the order of the pattern vectors as described in “Poll Method” on page 6-26.

See “Poll Options” on page 9-5 for more information.

## Search Options

Search options specify an optional search that the algorithm can perform at each iteration prior to the polling. If the search returns a point that improves the objective function, the algorithm uses that point at the next iteration and omits the polling. Please note, if you have selected the same **Search method** and **Poll method**, only the option selected in the Poll method will be used, although both will be used when the options selected are different.

**Complete search** (CompleteSearch) applies when you set **Search method** to GPS Positive basis Np1, GPS Positive basis 2N, MADS Positive basis Np1, MADS Positive basis 2N, or Latin hypercube. **Complete search** can have the values On or Off.

For GPS Positive basis Np1, MADS Positive basis Np1, GPS Positive basis 2N, or MADS Positive basis 2N, **Complete search** has the same meaning as the poll option **Complete poll**.

**Search method** (SearchMethod) specifies the optional search step. The options are

- None ([]) (the default) specifies no search step.
- GPS Positive basis Np1 ('GPSPositiveBasisNp1') performs a search step of a pattern search using the GPS Positive Basis Np1 option.

- GPS Positive basis 2N ('GPSPositiveBasis2N') performs a search step of a pattern search using the GPS Positive Basis 2N option.
- MADS Positive basis Np1 ('MADSPositiveBasisNp1') performs a search step of a pattern search using the MADS Positive Basis Np1 option.
- MADS Positive basis 2N ('MADSPositiveBasis2N') performs a search step of a pattern search using the MADS Positive Basis 2N option.
- Genetic Algorithm (@searchga) specifies a search using the genetic algorithm. If you select Genetic Algorithm, two other options appear:
  - **Iteration limit** — Positive integer specifying the number of iterations of the pattern search for which the genetic algorithm search is performed. The default for **Iteration limit** is 1.
  - **Options** — Options structure for the genetic algorithm, which you can set using `gaoptimset`.

To change the default values of **Iteration limit** and **Options** at the command line, use the syntax

```
options = psoptimset('SearchMethod',...
                    {@searchga,iterlim,optionsGA})
```

where `iterlim` is the value of **Iteration limit** and `optionsGA` is the genetic algorithm options structure.

- Latin hypercube (@searchlhs) specifies a Latin hypercube search. The way the search is performed depends on the setting for **Complete search**:
  - If you set **Complete search** to `On`, the algorithm polls all the points that are randomly generated at each iteration by the Latin hypercube search and chooses the one with the smallest objective function value.
  - If you set **Complete search** to `Off` (the default), the algorithm stops the poll as soon as it finds one of the randomly generated points whose objective function value is less than that of the current point, and chooses that point for the next iteration.

If you select Latin hypercube, two other options appear:

- **Iteration limit** — Positive integer specifying the number of iterations of the pattern search for which the Latin hypercube search is performed. The default for **Iteration limit** is 1.

- **Design level** — A positive integer specifying the design level. The number of points searched equals the **Design level** multiplied by the number of independent variables for the objective function. The default for **Design level** is 15.

To change the default values of **Iteration limit** and **Design level** at the command line, use the syntax

```
options=psoptimset('SearchMethod', {@searchlhs,iterlim,level})
```

where `iterlim` is the value of **Iteration limit** and `level` is the value of **Design level**.

- Nelder-Mead (@searchneldermead) specifies a search using `fminsearch`, which uses the Nelder-Mead algorithm. If you select Nelder-Mead, two other options appear:

- **Iteration limit** — Positive integer specifying the number of iterations of the pattern search for which the Nelder-Mead search is performed. The default for **Iteration limit** is 1.
- **Options** — Options structure for the function `fminsearch`, which you can create using the function `optimset`.

To change the default values of **Iteration limit** and **Options** at the command line, use the syntax

```
options=psoptimset('SearchMethod',...  
                  {@searchga,iterlim,optionsNM})
```

where `iterlim` is the value of **Iteration limit** and `optionsNM` is the options structure.

- Custom enables you to write your own search function. To specify the search function using the Pattern Search Tool,
  - Set **Search function** to Custom.
  - Set **Function name** to `@myfun`, where `myfun` is the name of your function.

If you are using `patternsearch`, set

```
options = psoptimset('SearchMethod', @myfun);
```



To see a template that you can use to write your own search function, enter

```
edit searchfcn_template
```

The following section describes the structure of the search function.

## Structure of the Search Function

Your search function must have the following calling syntax.

```
function [successSearch,xBest,fBest,funccount] =  
searchfcn_template(fun,x,A,b,Aeq,beq,lb,ub, ...  
    optimValues,options)
```

The search function has the following input arguments:

- `fun` — Objective function
- `x` — Current point
- `A,b` — Linear inequality constraints
- `Aeq,beq` — Linear equality constraints
- `lb,ub` — Lower and upper bound constraints
- `optimValues` — Structure that enables you to set search options. The structure contains the following fields:
  - `x` — Current point
  - `fval` — Objective function value at `x`
  - `iteration` — Current iteration number
  - `funccount` — Counter for user function evaluation
  - `scale` — Scale factor used to scale the design points
  - `problemtype` — Flag passed to the search routines, indicating whether the problem is 'unconstrained', 'boundconstraints', or 'linearconstraints'. This field is a sub-problem type for nonlinear constrained problems.
  - `meshsize` — Current mesh size used in search step
  - `method` — Method used in last iteration

- `options` — Pattern search options structure

The function has the following output arguments:

- `successSearch` — A Boolean identifier indicating whether the search is successful or not
- `xBest, fBest` — Best point and best function value found by search method

---

**Note** If you set **Search method** to Genetic algorithm or Nelder-Mead, we recommend that you leave **Iteration limit** set to the default value 1, because performing these searches more than once is not likely to improve results.

---

- `funccount` — Number of user function evaluation in search method

See “Using a Search Method” on page 6-32 for an example.

## Mesh Options

Mesh options control the mesh that the pattern search uses. The following options are available.

**Initial size** (`InitialMeshSize`) specifies the size of the initial mesh, which is the length of the shortest vector from the initial point to a mesh point. **Initial size** should be a positive scalar. The default is 1.0.

**Max size** (`MaxMeshSize`) specifies a maximum size for the mesh. When the maximum size is reached, the mesh size does not increase after a successful iteration. **Max size** must be a positive scalar, and is only used when the GPS algorithm is selected as the Poll or Search method. The default value is `Inf`.

**Accelerator** (`MeshAccelerator`) specifies whether the **Contraction factor** is multiplied by 0.5 after each unsuccessful iteration. **Accelerator** can have the values `On` or `Off`, the default.

**Rotate** (MeshRotate) specifies whether the mesh vectors are multiplied by -1 when the mesh size is less than a small value. **Rotate** is only applied when **Poll method** is set to GPS Positive basis Np1.

---

**Note** Changing the setting of **Rotate** has no effect on the poll when **Poll method** is set to GPS Positive basis 2N, MADS Positive basis 2N, or MADS Positive basis Np1.

---

**Scale** (ScaleMesh) specifies whether the algorithm scales the mesh points by multiplying the pattern vectors by constants. **Scale** can have the values Off or On (the default).

**Expansion factor** (MeshExpansion) specifies the factor by which the mesh size is increased after a successful poll. The default value is 2.0, which means that the size of the mesh is multiplied by 2.0 after a successful poll. **Expansion factor** must be a positive scalar and is only used when a GPS method is selected as the Poll or Search method.

**Contraction factor** (MeshContraction) specifies the factor by which the mesh size is decreased after an unsuccessful poll. The default value is 0.5, which means that the size of the mesh is multiplied by 0.5 after an unsuccessful poll. **Contraction factor** must be a positive scalar and is only used when a GPS method is selected as the Poll or Search method.

See “Mesh Expansion and Contraction” on page 6-35 for more information.

## Algorithm Settings

Algorithm settings define algorithmic specific parameters.

Parameters that can be specified for a nonlinear constraint algorithm include

- **Initial penalty** (InitialPenalty) — Specifies an initial value of the penalty parameter that is used by the algorithm. **Initial penalty** must be greater than or equal to 1.

- **Penalty factor** (`PenaltyFactor`) — Increases the penalty parameter when the problem is not solved to required accuracy and constraints are not satisfied. **Penalty factor** must be greater than 1.

## Cache Options

The pattern search algorithm can keep a record of the points it has already polled, so that it does not have to poll the same point more than once. If the objective function requires a relatively long time to compute, the cache option can speed up the algorithm. The memory allocated for recording the points is called the cache. This option should only be used for deterministic objective functions, but not for stochastic ones.

**Cache** (`Cache`) specifies whether a cache is used. The options are `On` and `Off`, the default. When you set **Cache** to `On`, the algorithm does not evaluate the objective function at any mesh points that are within **Tolerance** of a point in the cache.

**Tolerance** (`CacheTol`) specifies how close a mesh point must be to a point in the cache for the algorithm to omit polling it. **Tolerance** must be a positive scalar. The default value is `eps`.

**Size** (`CacheSize`) specifies the size of the cache. **Size** must be a positive scalar. The default value is `1e4`.

See “Using Cache” on page 6-41 for more information.

## Stopping Criteria

Stopping criteria determine what causes the pattern search algorithm to stop. Pattern search uses the following criteria:

**Mesh tolerance** (`TolMesh`) specifies the minimum tolerance for mesh size. The algorithm stops if the mesh size becomes smaller than **Mesh tolerance**. The default value is `1e-6`.

**Max iteration** (`MaxIter`) specifies the maximum number of iterations the algorithm performs. The algorithm stops if the number of iterations reaches **Max iteration**. You can select either

- **100\*numberofvariables** — Maximum number of iterations is 100 times the number of independent variables (the default).
- **Specify** — A positive integer for the maximum number of iterations

**Max function evaluations** (MaxFunEval) specifies the maximum number of evaluations of the objective function. The algorithm stops if the number of function evaluations reaches **Max function evaluations**. You can select either

- **2000\*numberofvariables** — Maximum number of function evaluations is 2000 times the number of independent variables.
- **Specify** — A positive integer for the maximum number of function evaluations

**Time limit** (TimeLimit) specifies the maximum time in seconds the pattern search algorithm runs before stopping. This also includes any specified pause time for pattern search algorithms.

**Bind tolerance** (TolBind) specifies the minimum tolerance for the distance from the current point to the boundary of the feasible region. Bind tolerance specifies when a linear constraint is active. It is not a stopping criterion. The default value is  $1e-3$ .

**X tolerance** (TolX) specifies the minimum distance between the current points at two consecutive iterations. The algorithm stops if the distance between two consecutive points is less than **X tolerance**. The default value is  $1e-6$ .

**Function tolerance** (TolFun) specifies the minimum tolerance for the objective function. After a successful poll, if the difference between the function value at the previous best point and function value at the current best point is less than the value of Function tolerance, the algorithm halts. The default value is  $1e-6$ .

See “Setting Tolerances for the Solver” on page 6-44 for an example.

**Nonlinear constraint tolerance** (TolCon) — The **Nonlinear constraint tolerance** is not used as stopping criterion. It is used to determine the feasibility with respect to nonlinear constraints.

## Output Function Options

Output functions are functions that the pattern search algorithm calls at each iteration. The following options are available:

- **History to new window** (@psoutputhistory) displays the history of points computed by the algorithm in the MATLAB Command Window at each multiple of **Interval** iterations.
- **Custom** enables you to write your own output function. To specify the output function using the Pattern Search Tool,
  - Select **Custom function**.
  - Enter @myfun in the text box, where myfun is the name of your function.

If you are using patternsearch, set

```
options = psoptimset('OutputFcn', @myfun);
```

To see a template that you can use to write your own output function, enter

```
edit psoutputfcn_template
```

at the MATLAB command prompt.

The following section describes the structure of the output function.

### Structure of the Output Function

Your output function must have the following calling syntax:

```
[stop, options, optchanged] =  
psoutputhistory(optimvalues, options, flag, interval)
```

The function has the following input arguments:

- **optimvalues** — Structure containing information about the current state of the solver. The structure contains the following fields:
  - **x** — Current point
  - **iteration** — Iteration number
  - **fval** — Objective function value

- `meshsize` — Current mesh size
- `funccount` — Number of function evaluations
- `method` — Method used in last iteration
- `TolFun` — Tolerance on function value in last iteration
- `TolX` — Tolerance on `x` value in last iteration
- `nonlinlineq` — Nonlinear inequality constraints, displayed only when a nonlinear constraint function is specified
- `nonlineq` — Nonlinear equality constraints, displayed only when a nonlinear constraint function is specified
- `options` — Options structure
- `flag` — Current state in which the output function is called. The possible values for `flag` are
  - `init` — Initialization state
  - `iter` — Iteration state
  - `interrupt` — Intermediate stage
  - `done` — Final state
- `interval` — Optional interval argument

“Parameterizing Functions Called by `patternsearch`” on page 6-53 explains how to provide additional parameters to the output function.

The output function returns the following arguments to `ga`:

- `stop` — Provides a way to stop the algorithm at the current iteration. `stop` can have the following values.
  - `false` — The algorithm continues to the next iteration.
  - `true` — The algorithm terminates at the current iteration.
- `options` — Options structure.
- `optchanged` — Flag indicating changes to options.

## Display to Command Window Options

**Level of display** ('Display') specifies how much information is displayed at the command line while the pattern search is running. The available options are

- `Off` ('off') — No output is displayed.
- `Iterative` ('iter') — Information is displayed for each iteration.
- `Diagnose` ('diagnose') — Information is displayed for each iteration. In addition, the diagnostic lists some problem information and the options that are changed from the defaults.
- `Final` ('final') — The reason for stopping is displayed.

Both `Iterative` and `Diagnose` display the following information:

- `Iter` — Iteration number
- `FunEval` — Cumulative number of function evaluations
- `MeshSize` — Current mesh size
- `FunVal` — Objective function value of the current point
- `Method` — Outcome of the current poll (with no nonlinear constraint function specified). With a nonlinear constraint function, `Method` displays the update method used after a subproblem is solved.
- `Max Constraint` — Maximum nonlinear constraint violation (displayed only when a nonlinear constraint function has been specified)

The default value of **Level of display** is

- `Off` in the Pattern Search Tool
- 'final' in an options structure created using `psoptimset`

## Vectorize Option

`Objective function is vectorized` specifies whether the computation of the objective function is vectorized. When **Objective function is vectorized** is `Off`, the algorithm calls the objective function on one point at a time as it loops through all of the mesh points. When **Objective function is vectorized** is `On`, the pattern search algorithm calls the objective function on all the points



in the mesh at once, i.e., in a single call to the objective function if either **Complete Poll** or **Complete Search** is On.

## Options Table for Pattern Search Algorithms

### Option Availability Table for GPS and MADS Algorithms

| Option       | Description  | Algorithm Availability |
|--------------|--|------------------------|
| Cache        | With Cache set to 'on', patternsearch keeps a history of the mesh points it polls and does not poll points close to them again at subsequent iterations. Use this option if patternsearch runs slowly because it is taking a long time to compute the objective function. If the objective function is stochastic, it is advised not to use this option. | GPS<br>MADS            |
| CacheSize    | Size of the cache.   | GPS<br>MADS            |
| CacheTol     | Positive scalar specifying how close the current mesh point must be to a point in the cache in order for patternsearch to avoid polling it. Available if 'Cache' option is set to 'on'.  | GPS<br>MADS            |
| CompletePoll | Complete poll around current iterate. Evaluate all the points in a poll step.  | GPS<br>MADS            |

**Option Availability Table for GPS and MADS Algorithms (Continued)**

| <b>Option</b>   | <b>Description</b>  | <b>Algorithm Availability</b>   |
|-----------------|---|---|
| CompleteSearch  | Complete poll around current iterate. Evaluate all the points in a search step. | GPS<br>MADS   |
| Display         | Level of display to Command Window.   | GPS<br>MADS   |
| InitialMeshSize | Initial mesh size used in pattern search algorithms.                            | GPS<br>MADS   |
| InitialPenalty  | Initial value of the penalty parameter.   | GPS<br>MADS   |
| MaxFunEvals     | Maximum number of objective function evaluations.                               | GPS<br>MADS   |
| MaxIter         | Maximum number of iterations.   | GPS<br>MADS   |
| MaxMeshSize     | Maximum mesh size used in a poll/search step.                                   | GPS<br>MADS   |
| MeshAccelerator | Accelerate mesh size contraction.   | GPS<br>MADS   |
| MeshContraction | Mesh contraction factor, used when iteration is unsuccessful.                   | GPS — Default value is 0.5.<br>MADS — Default value is 0.25 if MADS algorithm is selected for either the <b>Poll method</b> or <b>Search method</b> . |

**Option Availability Table for GPS and MADS Algorithms (Continued)**

| <b>Option</b> | <b>Description</b>   | <b>Algorithm Availability</b>  |
|---------------|--|--|
| MeshExpansion | Mesh expansion factor, expands mesh when iteration is successful.      | GPS — Default value is 2.<br>MADS — Default value is 4 if MADS algorithm is selected for either the <b>Poll method</b> or <b>Search method</b> . |
| MeshRotate    | Rotate the pattern before declaring a point to be optimum.             | GPS<br>MADS  |
| OutputFcn     | User-specified function that a pattern search calls at each iteration. | GPS<br>MADS  |
| PenaltyFactor | Penalty update parameter.  | GPS<br>MADS  |
| PlotFcn       | Specifies function to plot at runtime.                                 | GPS<br>MADS  |
| PlotInterval  | Specifies that plot functions will be called at every interval.        | GPS<br>MADS  |
| PollingOrder  | Order in which search directions are polled.                           | GPS only   |
| PollMethod    | Polling strategy used in pattern search.                               | GPS<br>MADS  |
| ScaleMesh     | Automatic scaling of variables.  | GPS<br>MADS  |
| SearchMethod  | Specifies search method used in pattern search.                        | GPS<br>MADS  |

**Option Availability Table for GPS and MADS Algorithms (Continued)**

| <b>Option</b> | <b>Description</b>  | <b>Algorithm Availability</b> |
|---------------|---|-------------------------------|
| TimeLimit     | Total time (in seconds) allowed for optimization. Also includes any specified pause time for pattern search algorithms. | GPS<br>MADS                   |
| TolBind       | Binding tolerance used to determine if linear constraint is active.   | GPS<br>MADS                   |
| TolCon        | Tolerance on nonlinear constraints.   | GPS<br>MADS                   |
| TolFun        | Tolerance on function value.  | GPS<br>MADS                   |
| TolMesh       | Tolerance on mesh size.   | GPS<br>MADS                   |
| TolX          | Tolerance on independent variable.  | GPS<br>MADS                   |
| Vectorized    | Specifies whether functions are vectorized.   | GPS<br>MADS                   |

## Genetic Algorithm Options

This section describes the options for the genetic algorithm. There are two ways to specify options, depending on whether you are using the Genetic Algorithm Tool or calling the function `ga` at the command line:

- If you are using the Genetic Algorithm Tool (`gatool`), you specify the options by selecting an option from a drop-down list or by entering the value of the option in a text field. See “Setting Options in the Genetic Algorithm Tool” on page 7-14.
- If you are calling `ga` from the command line, you specify the options by creating an options structure using the function `gaoptimset`, as follows:

```
options = gaoptimset('Param1', value1, 'Param2', value2, ...);
```

See “Setting Options for `ga` at the Command Line” on page 7-26 for examples.

In this section, each option is listed in two ways:

- By its label, as it appears in the Genetic Algorithm Tool
- By its field name in the options structure

For example:

- **Population type** refers to the label of the option in the Genetic Algorithm Tool.
- `PopulationType` refers to the corresponding field of the `options` structure.

The genetic algorithm options are divided into the following categories:

- “Plot Options” on page 9-24
- “Population Options” on page 9-27
- “Fitness Scaling Options” on page 9-29
- “Selection Options” on page 9-31
- “Reproduction Options” on page 9-32
- “Mutation Options” on page 9-33

- “Crossover Options” on page 9-35
- “Migration Options” on page 9-39
- “Algorithm Settings” on page 9-40
- “Hybrid Function Options” on page 9-40
- “Stopping Criteria Options” on page 9-40
- “Output Function Options” on page 9-41
- “Display to Command Window Options” on page 9-43
- “Vectorize Option” on page 9-44

## Plot Options

Plot options enable you to plot data from the genetic algorithm while it is running. When you select plot functions and run the genetic algorithm, a plot window displays the plots on separate axes. Click on any subplot to view a larger version of the plot in a separate figure window. You can stop the algorithm at any time by clicking the **Stop** button on the plot window.

**Plot interval** (`PlotInterval`) specifies the number of generations between consecutive calls to the plot function.

You can select any of the following plot functions in the **Plots** pane:

- **Best fitness** (`@gaplotbestf`) plots the best function value versus generation.
- **Expectation** (`@gaplotexpectation`) plots the expected number of children versus the raw scores at each generation.
- **Score diversity** (`@gaplotscorediversity`) plots a histogram of the scores at each generation.
- **Stopping** (`@plotstopping`) plots stopping criteria levels.
- **Best individual** (`@gaplotbestindiv`) plots the vector entries of the individual with the best fitness function value in each generation.
- **Genealogy** (`@gaplotgenealogy`) plots the genealogy of individuals. Lines from one generation to the next are color-coded as follows:
  - Red lines indicate mutation children.

- Blue lines indicate crossover children.
  - Black lines indicate elite individuals.
  - **Scores** (@gaplotscores) plots the scores of the individuals at each generation.
  - **Max constraint** (@gaplotmaxconstr) plots the maximum nonlinear constraint violation at each generation.
  - **Distance** (@gaplotdistance) plots the average distance between individuals at each generation.
  - **Range** (@gaplotrange) plots the minimum, maximum, and mean fitness function values in each generation.
  - **Selection** (@gaplotselection) plots a histogram of the parents.
  - **Custom function** enables you to use plot functions of your own. To specify the plot function if you are using the Genetic Algorithm Tool,
    - Select **Custom function**.
    - Enter @myfun in the text box, where myfun is the name of your function.
- See “Structure of the Plot Functions” on page 9-25.

To display a plot when calling `ga` from the command line, set the `PlotFcns` field of options to be a function handle to the plot function. For example, to display the best fitness plot, set options as follows

```
options = gaoptimset('PlotFcns', @gaplotbestf);
```

To display multiple plots, use the syntax

```
options =gaoptimset('PlotFcns', {@plotfun1, @plotfun2, ...});
```

where @plotfun1, @plotfun2, and so on are function handles to the plot functions.

### Structure of the Plot Functions

The first line of a plot function has the form

```
function state = plotfun(options, state, flag)
```

The input arguments to the function are

- `options` — Structure containing all the current options settings.
- `state` — Structure containing information about the current generation. “The State Structure” on page 9-26 describes the fields of state.
- `flag` — String that tells what stage the algorithm is currently in.

“Parameterizing Functions Called by `patternsearch`” on page 6-53 explains how to provide additional parameters to the function.

### **The State Structure**

The state structure, which is an input argument to `plot`, `mutation`, and `output` functions, contains the following fields:

- `Population` — Population in the current generation
- `Score` — Scores of the current population
- `Generation` — Current generation number
- `StartTime` — Time when genetic algorithm started
- `StopFlag` — String containing the reason for stopping
- `Selection` — Indices of individuals selected for elite, crossover and mutation
- `Expectation` — Expectation for selection of individuals
- `Best` — Vector containing the best score in each generation
- `LastImprovement` — Generation at which the last improvement in fitness value occurred
- `LastImprovementTime` — Time at which last improvement occurred
- `NonlinIneq` — Nonlinear inequality constraints, displayed only when a nonlinear constraint function is specified
- `NonlinEq` — Nonlinear equality constraints, displayed only when a nonlinear constraint function is specified



## Population Options

Population options enable you to specify the parameters of the population that the genetic algorithm uses.

**Population type** (`PopulationType`) specifies the data type of the input to the fitness function. You can set **Population type** to be one of the following:

- `Double Vector ('doubleVector')` — Use this option if the individuals in the population have type `double`. This is the default.
- `Bit string ('bitstring')` — Use this option if the individuals in the population are bit strings.
- `Custom ('custom')` — Use this option to create a population whose data type is neither of the preceding.

If you use a custom population type, you must write your own creation, mutation, and crossover functions that accept inputs of that population type, and specify these functions in the following fields, respectively:

- **Creation function** (`CreationFcn`)
- **Mutation function** (`MutationFcn`)
- **Crossover function** (`CrossoverFcn`)

**Population size** (`PopulationSize`) specifies how many individuals there are in each generation. With a large population size, the genetic algorithm searches the solution space more thoroughly, thereby reducing the chance that the algorithm will return a local minimum that is not a global minimum. However, a large population size also causes the algorithm to run more slowly.

If you set **Population size** to a vector, the genetic algorithm creates multiple subpopulations, the number of which is the length of the vector. The size of each subpopulation is the corresponding entry of the vector.

**Creation function** (`CreationFcn`) specifies the function that creates the initial population for `ga`. You can choose from the following functions:

- `Uniform (@gacreationuniform)` creates a random initial population with a uniform distribution. This is the default.

- Custom enables you to write your own creation function, which must generate data of the type that you specify in **Population type**. To specify the creation function if you are using the Genetic Algorithm Tool,
  - Set **Creation function** to Custom.
  - Set **Function name** to @myfun, where myfun is the name of your function.

If you are using ga, set

```
options = gaoptimset('CreationFcn', @myfun);
```

Your creation function must have the following calling syntax.

```
function Population = myfun(GenomeLength, FitnessFcn, options)
```

The input arguments to the function are

- GenomeLength — Number of independent variables for the fitness function
- FitnessFcn — Fitness function
- options — Options structure

The function returns Population, the initial population for the genetic algorithm.

“Parameterizing Functions Called by patternsearch” on page 6-53 explains how to provide additional parameters to the function.

**Initial population** (InitialPopulation) specifies an initial population for the genetic algorithm. The default value is [], in which case ga uses the **Creation function** to create an initial population. If you enter a nonempty array in the **Initial population** field, the array must have **Population size** rows and **Number of variables** columns. In this case, the genetic algorithm does not call the **Creation function**.

**Initial scores** (InitialScores) specifies initial scores for the initial population.

**Initial range** (PopInitRange) specifies the range of the vectors in the initial population that is generated by the creation function. You can set **Initial**

**range** to be a matrix with two rows and **Number of variables** columns, each column of which has the form [lb; ub], where lb is the lower bound and ub is the upper bound for the entries in that coordinate. If you specify **Initial range** to be a 2-by-1 vector, each entry is expanded to a constant row of length **Number of variables**.

See “Example — Setting the Initial Range” on page 7-35 for an example.

## Fitness Scaling Options

Fitness scaling converts the raw fitness scores that are returned by the fitness function to values in a range that is suitable for the selection function. You can specify options for fitness scaling in the **Fitness scaling** pane.

**Scaling function** (FitnessScalingFcn) specifies the function that performs the scaling. The options are

- Rank (@fitscalingrank) — The default fitness scaling function, Rank, scales the raw scores based on the rank of each individual instead of its score. The rank of an individual is its position in the sorted scores. The rank of the most fit individual is 1, the next most fit is 2, and so on. Rank fitness scaling removes the effect of the spread of the raw scores.
- Proportional (@fitscalingprop) — Proportional scaling makes the scaled value of an individual proportional to its raw fitness score.
- Top (@fitscalingtop) — Top scaling scales the top individuals equally. Selecting Top displays an additional field, **Quantity**, which specifies the number of individuals that are assigned positive scaled values. **Quantity** can be an integer between 1 and the population size or a fraction between 0 and 1 specifying a fraction of the population size. The default value is 0.4. Each of the individuals that produce offspring is assigned an equal scaled value, while the rest are assigned the value 0. The scaled values have the form [0 1/n 1/n 0 0 1/n 0 0 1/n ...].

To change the default value for **Quantity** at the command line, use the following syntax

```
options = gaoptimset('FitnessScalingFcn', {@fitscalingtop,
quantity})
```

where quantity is the value of **Quantity**.

- `Shift linear (@fitscalingshiftlinear)` — Shift linear scaling scales the raw scores so that the expectation of the fittest individual is equal to a constant multiplied by the average score. You specify the constant in the **Max survival rate** field, which is displayed when you select Shift linear. The default value is 2.

To change the default value of **Max survival rate** at the command line, use the following syntax

```
options = gaoptimset('FitnessScalingFcn',  
    {@fitscalingshiftlinear, rate})
```

where `rate` is the value of **Max survival rate**.

- `Custom` enables you to write your own scaling function. To specify the scaling function using the Genetic Algorithm Tool,
  - Set **Scaling function** to Custom.
  - Set **Function name** to `@myfun`, where `myfun` is the name of your function.

If you are using `ga` at the command line, set

```
options = gaoptimset('FitnessScalingFcn', @myfun);
```

Your scaling function must have the following calling syntax:

```
function expectation = myfun(scores, nParents)
```

The input arguments to the function are

- `scores` — A vector of scalars, one for each member of the population
- `nParents` — The number of parents needed from this population

The function returns `expectation`, a row vector of scalars of the same length as `scores`, giving the scaled values of each member of the population. The sum of the entries of `expectation` must equal `nParents`.

“Parameterizing Functions Called by `patternsearch`” on page 6-53 explains how to provide additional parameters to the function.

See “Fitness Scaling” on page 7-39 for more information.

## Selection Options

Selection options specify how the genetic algorithm chooses parents for the next generation. You can specify the function the algorithm uses in the **Selection function** (`SelectionFcn`) field in the **Selection** options pane. The options are

- **Stochastic uniform** (`@selectionstochunif`) — The default selection function, **Stochastic uniform**, lays out a line in which each parent corresponds to a section of the line of length proportional to its scaled value. The algorithm moves along the line in steps of equal size. At each step, the algorithm allocates a parent from the section it lands on. The first step is a uniform random number less than the step size.
- **Remainder** (`@selectionremainder`) — **Remainder selection** assigns parents deterministically from the integer part of each individual's scaled value and then uses roulette selection on the remaining fractional part. For example, if the scaled value of an individual is 2.3, that individual is listed twice as a parent because the integer part is 2. After parents have been assigned according to the integer parts of the scaled values, the rest of the parents are chosen stochastically. The probability that a parent is chosen in this step is proportional to the fractional part of its scaled value.
- **Uniform** (`@selectionuniform`) — **Uniform selection** chooses parents using the expectations and number of parents. **Uniform selection** is useful for debugging and testing, but is not a very effective search strategy.
- **Roulette** (`@selectionroulette`) — **Roulette selection** chooses parents by simulating a roulette wheel, in which the area of the section of the wheel corresponding to an individual is proportional to the individual's expectation. The algorithm uses a random number to select one of the sections with a probability equal to its area.
- **Tournament** (`@selectiontournament`) — **Tournament selection** chooses each parent by choosing **Tournament size** players at random and then choosing the best individual out of that set to be a parent. **Tournament size** must be at least 2. The default value of **Tournament size** is 4.

To change the default value of **Tournament size** at the command line, use the syntax

```
options = gaoptimset('SelectionFcn',...
                    {@selecttournament,size})
```

where `size` is the value of **Tournament size**.

- **Custom** enables you to write your own selection function. To specify the selection function using the Genetic Algorithm Tool,
  - Set **Selection function** to **Custom**.
  - Set **Function name** to `@myfun`, where `myfun` is the name of your function.

If you are using `ga` at the command line, set

```
options = gaoptimset('SelectionFcn', @myfun);
```

Your selection function must have the following calling syntax:

```
function parents = myfun(expectation, nParents, options)
```

The input arguments to the function are

- `expectation` — Expected number of children for each member of the population
- `nParents` — Number of parents to select
- `options` — Genetic algorithm options structure

The function returns `parents`, a row vector of length `nParents` containing the indices of the parents that you select.

“Parameterizing Functions Called by `patternsearch`” on page 6-53 explains how to provide additional parameters to the function.

See “Selection” on page 7-42 for more information.

## Reproduction Options

Reproduction options specify how the genetic algorithm creates children for the next generation.

**Elite count** (`EliteCount`) specifies the number of individuals that are guaranteed to survive to the next generation. Set **Elite count** to be a positive integer less than or equal to the population size. The default value is 2.

**Crossover fraction** (CrossoverFraction) specifies the fraction of the next generation, other than elite children, that are produced by crossover. Set **Crossover fraction** to be a fraction between 0 and 1, either by entering the fraction in the text box or moving the slider. The default value is 0.8.

See “Setting the Crossover Fraction” on page 7-46 for an example.

## Mutation Options

Mutation options specify how the genetic algorithm makes small random changes in the individuals in the population to create mutation children. Mutation provides genetic diversity and enable the genetic algorithm to search a broader space. You can specify the mutation function in the **Mutation function** (MutationFcn) field in the **Mutation** options pane. You can choose from the following functions:

- **Gaussian** (mutationgaussian) — The default mutation function, **Gaussian**, adds a random number taken from a Gaussian distribution with mean 0 to each entry of the parent vector. The variance of this distribution is determined by the parameters **Scale** and **Shrink**, which are displayed when you select **Gaussian**, and by the **Initial range** setting in the **Population** options.
  - The **Scale** parameter determines the variance at the first generation. If you set **Initial range** to be a 2-by-1 vector  $v$ , the initial variance is the same at all coordinates of the parent vector, and is given by  $\mathbf{Scale} * (v(2) - v(1))$ .  
If you set **Initial range** to be a vector  $v$  with two rows and **Number of variables** columns, the initial variance at coordinate  $i$  of the parent vector is given by  $\mathbf{Scale} * (v(i,2) - v(i,1))$ .
  - The **Shrink** parameter controls how the variance shrinks as generations go by. If you set **Initial range** to be a 2-by-1 vector, the variance at the  $k$ th generation,  $var_k$ , is the same at all coordinates of the parent vector, and is given by the recursive formula

$$var_k = var_{k-1} \left( 1 - \mathbf{Shrink} \cdot \frac{k}{\mathbf{Generations}} \right)$$

If you set **Initial range** to be a vector with two rows and **Number of variables** columns, the variance at coordinate  $i$  of the parent vector at the  $k$ th generation,  $var_{i,k}$ , is given by the recursive formula

$$var_{i,k} = var_{i,k-1} \left( 1 - \text{Shrink} \cdot \frac{k}{\text{Generations}} \right)$$

If you set **Shrink** to 1, the algorithm shrinks the variance in each coordinate linearly until it reaches 0 at the last generation is reached. A negative value of **Shrink** causes the variance to grow.

The default values of **Scale** and **Shrink** are 0.5 and 0.75, respectively. To change these default values at the command line, use the syntax

```
options = gaoptimset('MutationFcn', ...
    {@mutationgaussian, scale, shrink})
```

where `scale` and `shrink` are the values of **Scale** and **Shrink**, respectively.

- **Uniform** (`mutationuniform`) — Uniform mutation is a two-step process. First, the algorithm selects a fraction of the vector entries of an individual for mutation, where each entry has a probability **Rate** of being mutated. The default value of **Rate** is 0.01. In the second step, the algorithm replaces each selected entry by a random number selected uniformly from the range for that entry.

To change the default value of **Rate** at the command line, use the syntax

```
options = gaoptimset('MutationFcn', {@mutationuniform, rate})
```

where `rate` is the value of **Rate**.

- **Adaptive Feasible** (`mutationadaptfeasible`) randomly generates directions that are adaptive with respect to the last successful or unsuccessful generation. The feasible region is bounded by the constraints and inequality constraints. A step length is chosen along each direction so that linear constraints and bounds are satisfied.
- **Custom** enables you to write your own mutation function. To specify the mutation function using the Genetic Algorithm Tool,
  - Set **Mutation function** to Custom.



- Set **Function name** to @myfun, where myfun is the name of your function.

If you are using ga, set

```
options = gaoptimset('MutationFcn', @myfun);
```

Your mutation function must have this calling syntax:

```
function mutationChildren = myfun(parents, options, nvars,
    FitnessFcn, state, thisScore, thisPopulation)
```

The arguments to the function are

- parents — Row vector of parents chosen by the selection function
- options — Options structure
- nvars — Number of variables
- FitnessFcn — Fitness function
- state — Structure containing information about the current generation. “The State Structure” on page 9-26 describes the fields of state.
- thisScore — Vector of scores of the current population
- thisPopulation — Matrix of individuals in the current population

The function returns mutationChildren—the mutated offspring—as a matrix whose rows correspond to the children. The number of columns of the matrix is **Number of variables**.

“Parameterizing Functions Called by patternsearch” on page 6-53 explains how to provide additional parameters to the function.

## Crossover Options

Crossover options specify how the genetic algorithm combines two individuals, or parents, to form a crossover child for the next generation.

**Crossover function** (CrossoverFcn) specifies the function that performs the crossover. You can choose from the following functions:

- Scattered (@crossoverscattered), the default crossover function, creates a random binary vector and selects the genes where the vector is a 1 from

the first parent, and the genes where the vector is a 0 from the second parent, and combines the genes to form the child. For example, if p1 and p2 are the parents

```
p1 = [a b c d e f g h]
p2 = [1 2 3 4 5 6 7 8]
```

and the binary vector is [1 1 0 0 1 0 0 0], the function returns the following child:

```
child1 = [a b 3 4 e 6 7 8]
```

- Single point (@crossoversinglepoint) chooses a random integer n between 1 and **Number of variables** and then
  - Selects vector entries numbered less than or equal to n from the first parent.
  - Selects vector entries numbered greater than n from the second parent.
  - Concatenates these entries to form a child vector.

For example, if p1 and p2 are the parents

```
p1 = [a b c d e f g h]
p2 = [1 2 3 4 5 6 7 8]
```

and the crossover point is 3, the function returns the following child.

```
child = [a b c 4 5 6 7 8]
```

- Two point (@crossovertwopoint) selects two random integers m and n between 1 and **Number of variables**. The function selects
  - Vector entries numbered less than or equal to m from the first parent
  - Vector entries numbered from m+1 to n, inclusive, from the second parent
  - Vector entries numbered greater than n from the first parent.

The algorithm then concatenates these genes to form a single gene. For example, if p1 and p2 are the parents

```
p1 = [a b c d e f g h]
p2 = [1 2 3 4 5 6 7 8]
```

and the crossover points are 3 and 6, the function returns the following child.

```
child = [a b c 4 5 6 g h]
```

- **Intermediate** (@crossoverintermediate) creates children by taking a weighted average of the parents. You can specify the weights by a single parameter, **Ratio**, which can be a scalar or a row vector of length **Number of variables**. The default is a vector of all 1's. The function creates the child from parent1 and parent2 using the following formula.

$$\text{child} = \text{parent1} + \text{rand} * \mathbf{Ratio} * (\text{parent2} - \text{parent1})$$

If all the entries of **Ratio** lie in the range [0, 1], the children produced are within the hypercube defined by placing the parents at opposite vertices. If **Ratio** is not in that range, the children might lie outside the hypercube. If **Ratio** is a scalar, then all the children lie on the line between the parents.

To change the default value of **Ratio** at the command line, use the syntax

```
options = gaoptimset('CrossoverFcn', ...
    {@crossoverintermediate, ratio});
```

where `ratio` is the value of **Ratio**.

- **Heuristic** (@crossoverheuristic) returns a child that lies on the line containing the two parents, a small distance away from the parent with the better fitness value in the direction away from the parent with the worse fitness value. You can specify how far the child is from the better parent by the parameter **Ratio**, which appears when you select **Heuristic**. The default value of **Ratio** is 1.2. If parent1 and parent2 are the parents, and parent1 has the better fitness value, the function returns the child

$$\text{child} = \text{parent2} + R * (\text{parent1} - \text{parent2});$$

To change the default value of **Ratio** at the command line, use the syntax

```
options=gaoptimset('CrossoverFcn',...
    {@crossoverheuristic,ratio});
```

where `ratio` is the value of **Ratio**.

- Arithmetic (@crossoverarithmetic) creates children that are the weighted arithmetic mean of two parents. Children are always feasible with respect to linear constraints and bounds.
- Custom enables you to write your own crossover function. To specify the crossover function using the Genetic Algorithm Tool,
  - Set **Crossover function** to Custom.
  - Set **Function name** to @myfun, where myfun is the name of your function.

If you are using ga, set

```
options = gaoptimset('CrossoverFcn', @myfun);
```

Your selection function must have the following calling syntax.

```
xoverKids = myfun(parents, options, nvars, FitnessFcn,  
unused, thisPopulation)
```

The arguments to the function are

- parents — Row vector of parents chosen by the selection function
- options — options structure
- nvars — Number of variables
- FitnessFcn — Fitness function
- unused — Placeholder not used
- thisPopulation — Matrix representing the current population. The number of rows of the matrix is **Population size** and the number of columns is **Number of variables**.

The function returns xoverKids—the crossover offspring—as a matrix whose rows correspond to the children. The number of columns of the matrix is **Number of variables**.

“Parameterizing Functions Called by patternsearch” on page 6-53 explains how to provide additional parameters to the function.

## Migration Options

Migration options specify how individuals move between subpopulations. Migration occurs if you set **Population size** to be a vector of length greater than 1. When migration occurs, the best individuals from one subpopulation replace the worst individuals in another subpopulation. Individuals that migrate from one subpopulation to another are copied. They are not removed from the source subpopulation.

You can control how migration occurs by the following three fields in the **Migration** options pane:

- **Direction** (MigrationDirection) — Migration can take place in one or both directions.
  - If you set **Direction** to Forward ('forward'), migration takes place toward the last subpopulation. That is, the *n*th subpopulation migrates into the (*n*+1)th subpopulation.
  - If you set **Direction** to Both ('both'), the *n*th subpopulation migrates into both the (*n*-1)th and the (*n*+1)th subpopulation.

Migration wraps at the ends of the subpopulations. That is, the last subpopulation migrates into the first, and the first may migrate into the last.

- **Interval** (MigrationInterval) — Specifies how many generation pass between migrations. For example, if you set **Interval** to 20, migration takes place every 20 generations.
- **Fraction** (MigrationFraction) — Specifies how many individuals move between subpopulations. **Fraction** specifies the fraction of the smaller of the two subpopulations that moves. For example, if individuals migrate from a subpopulation of 50 individuals into a subpopulation of 100 individuals and you set **Fraction** to 0.1, the number of individuals that migrate is  $0.1 * 50 = 5$ .

## Algorithm Settings

Algorithm settings define algorithmic specific parameters.

Parameters that can be specified for a nonlinear constraint algorithm include

- **Initial penalty** (`InitialPenalty`) — Specifies an initial value of the penalty parameter that is used by the algorithm. **Initial penalty** must be greater than or equal to 1.
- **Penalty factor** (`PenaltyFactor`) — Increases the penalty parameter when the problem is not solved to required accuracy and constraints are not satisfied. **Penalty factor** must be greater than 1.

## Hybrid Function Options

A hybrid function is another minimization function that runs after the genetic algorithm terminates. You can specify a hybrid function in **Hybrid function** (`HybridFcn`) options. The choices are

- `[]` — No hybrid function
- `fminsearch` (`@fminsearch`) — Uses the MATLAB function `fminsearch` to perform unconstrained minimization
- `patternsearch` (`@patternsearch`) — Uses a pattern search to perform constrained or unconstrained minimization
- `fminunc` (`@fminunc`) — Uses the Optimization Toolbox function `fminunc` to perform unconstrained minimization
- `fmincon` (`@fmincon`) — Uses the Optimization Toolbox function `fmincon` to perform constrained minimization

See “Using a Hybrid Function” on page 7-56 for an example.

## Stopping Criteria Options

Stopping criteria determine what causes the algorithm to terminate. You can specify the following options:

- **Generations** (`Generations`) — Specifies the maximum number of iterations the genetic algorithm will perform. The default is 100.

- **Time limit** (TimeLimit) — Specifies the maximum time in seconds the genetic algorithm runs before stopping.
- **Fitness limit** (FitnessLimit) — The algorithm stops if the best fitness value is less than or equal to the value of **Fitness limit**.
- **Stall generations** (StallGenLimit) — The algorithm stops if the weighted average change in the fitness function value over **Stall generations** is less than **Function tolerance**.
- **Stall time limit** (StallTimeLimit) — The algorithm stops if there is no improvement in the best fitness value for an interval of time in seconds specified by **Stall time**.
- **Function tolerance** (TolFun) — The algorithm runs until the cumulative change in the fitness function value over **Stall generations** is less than **Function Tolerance**.
- **Nonlinear constraint tolerance** (TolCon) — The **Nonlinear constraint tolerance** is not used as stopping criterion. It is used to determine the feasibility with respect to nonlinear constraints.

See “Setting the Maximum Number of Generations” on page 7-59 for an example.

## Output Function Options

Output functions are functions that the genetic algorithm calls at each generation. The following options are available:

**History to new window** (@gaoutputgen) displays the history of points computed by the algorithm in a new window at each multiple of **Interval** iterations.

**Custom** enables you to write your own output function. To specify the output function using the Genetic Algorithm Tool,

- Select **Custom function**.
- Enter @myfun in the text box, where myfun is the name of your function.

If you are using ga, set

```
options = gaoptimset('OutputFcn', @myfun);
```

To see a template that you can use to write your own output functions, enter

```
edit gaoutputfcntemplate
```

at the MATLAB command line.

For information on output function options available with pattern search, see “Output Function Options” on page 9-16 in Pattern Search Options.

### Structure of the Output Function

The output function has the following calling syntax.

```
[state, options,optchanged] = myfun(options,state,flag,interval)
```

The function has the following input arguments:

- `options` — Options structure
- `state` — Structure containing information about the current generation. “The State Structure” on page 9-26 describes the fields of state.
- `flag` — String indicating the current status of the algorithm as follows:
  - `'init'` — Initial stage
  - `'iter'` — Algorithm running
  - `'interrupt'` — Intermediate stage
  - `'done'` — Algorithm terminated
- `interval` — Optional interval argument

“Parameterizing Functions Called by `patternsearch`” on page 6-53 explains how to provide additional parameters to the function.

The output function returns the following arguments to `ga`:

- `state` — Structure containing information about the current generation



- `options` — Options structure modified by the output function. This argument is optional.
- `optchanged` — Flag indicating changes to options

## Display to Command Window Options

**Level of display** ('Display') specifies how much information is displayed at the command line while the genetic algorithm is running. The available options are

- `Off` ('off') — No output is displayed.
- `Iterative` ('iter') — Information is displayed at each iteration.
- `Diagnose` ('diagnose') — Information is displayed at each iteration. In addition, the diagnostic lists some problem information and the options that have been changed from the defaults.
- `Final` ('final') — The reason for stopping is displayed.

Both `Iterative` and `Diagnose` display the following information:

- `Generation` — Generation number
- `f-count` — Cumulative number of fitness function evaluations
- `Best f(x)` — Best fitness function value
- `Mean f(x)` — Mean fitness function value
- `Stall generations` — Number of generations since the last improvement of the fitness function

When a nonlinear constraint function has been specified, `Iterative` and `Diagnose` will not display the `Mean f(x)`, but will additionally display:

- `Max Constraint` — Maximum nonlinear constraint violation

The default value of **Level of display** is

- `Off` in the Genetic Algorithm Tool
- `'final'` in an options structure created using `gaoptimset`

## Vectorize Option

The vectorize option specifies whether the computation of the fitness function is vectorized. Set **Set Objective function is vectorized** to **On** to indicate that the fitness function is vectorized. When **Objective function is vectorized** is **Off**, the algorithm calls the fitness function on one individual at a time as it loops through the population.

See “Vectorizing the Fitness Function” on page 7-60 for an example.

# Simulated Annealing and Threshold Acceptance Algorithm Options

This section describes the options for the simulated annealing and threshold acceptance solvers. You specify the options by creating an options structure using the function `saoptimset`, as follows:

```
options = saoptimset('Param1', value1, 'Param2', value2, ...);
```

See “Setting Options for `simulannealbnd` and `threshacceptbnd` at the Command Line” on page 8-3 for examples.

Each option in this section is listed by its field name in the options structure. For example, `InitialTemperature` refers to the corresponding field of the options structure.

The options for the simulated annealing and threshold acceptance solvers are divided into the following categories:

- “Plot Options” on page 9-45
- “Temperature Options” on page 9-48
- “Algorithm Settings” on page 9-49
- “Hybrid Function Options” on page 9-49
- “Stopping Criteria Options” on page 9-50
- “Output Function Options” on page 9-50
- “Display Options” on page 9-52

## Plot Options

Plot options enable you to plot data from the simulated annealing or threshold acceptance solver while it is running. When you specify plot functions and run the algorithm, a plot window displays the plots on separate axes. Right-click on any subplot to view a larger version of the plot in a separate figure window.

`PlotInterval` specifies the number of iterations between consecutive calls to the plot function.

To display a plot when calling `simulannealbnd` or `threshacceptbnd` from the command line, set the `PlotFcns` field of options to be a function handle to the plot function. You can specify any of the following plots:

- `@saplotbestf` plots the best objective function value.
- `@saplotbestx` plots the current best point.
- `@saplotf` plots the current function value.
- `@saplotx` plots the current point.
- `@saplotstopping` plots stopping criteria levels.
- `@saplottemperature` plots the temperature at each iteration.
- `@myfun` plots a custom plot function, where `myfun` is the name of your function. See “Structure of the Plot Functions” on page 9-4 for a description of the syntax.

For example, to display the best objective plot, set options as follows

```
options = saoptimset('PlotFcns', @saplotbestf);
```

To display multiple plots, use the cell array syntax

```
options = saoptimset('PlotFcns', {@plotfun1, @plotfun2, ...});
```

where `@plotfun1`, `@plotfun2`, and so on are function handles to the plot functions.

### Structure of the Plot Functions

The first line of a plot function has the form

```
function stop = plotfun(options, optimvalues, flag)
```

The input arguments to the function are

- `options` — Options structure created using `saoptimset`.
- `optimvalues` — Structure containing information about the current state of the solver. The structure contains the following fields:
  - `x` — Current point

- `fval` — Objective function value at `x`
- `bestx` — Best point found so far
- `bestfval` — Objective function value at best point
- `temperature` — Current temperature
- `iteration` — Current iteration
- `funccount` — Number of function evaluations
- `t0` — Start time for algorithm
- `k` — Annealing parameter
- `tau` — Threshold acceptance sequence (for threshold acceptance solver only)
- `flag` — Current state in which the plot function is called. The possible values for `flag` are
  - `'init'` — Initialization state
  - `'iter'` — Iteration state
  - `'done'` — Final state

The output argument `stop` provides a way to stop the algorithm at the current iteration. `stop` can have the following values:

- `false` — The algorithm continues to the next iteration.
- `true` — The algorithm terminates at the current iteration.

## Temperature Options

Temperature options specify how the temperature will be lowered at each iteration over the course of the algorithm.

- `InitialTemperature` — Initial temperature at the start of the algorithm. The default is 100.
- `TemperatureFcn` — Function used to update the temperature schedule. The choices are:
  - `@temperatureexp` — The temperature at an iteration is reduced exponentially at each iteration. This is the default.
  - `@temperaturefast` — Uses the fast annealing function.
  - `@temperaturebolz` — Uses the Boltzmann annealing function.
  - `@myfun` — Uses a custom function, `myfun`, to update temperature. See the functions above for syntax.
- `ReannealInterval` — Number of points accepted before reannealing. The default value is 100.

## Algorithm Settings

Algorithm settings define algorithmic specific parameters used in generating new points at each iteration.

Parameters that can be specified for the simulated annealing and threshold acceptance algorithms are:

- `AnnealingFcn` — Function used to generate new points for the next iteration. The choices are:
  - `@annealingfast` — Uses the fast annealing algorithm. This is the default.
  - `@annealingbolz` — Uses the Boltzmann annealing algorithm.
  - `@myfun` — Uses a custom annealing algorithm, `myfun`. See the functions above for syntax.

`AcceptanceFcn` — Function used to determine whether a new point is accepted or not. The default function varies depending on which solver/algorithm you are using. The choices are:

- `@acceptancesa` — Simulated annealing acceptance function. The default for `simulannealbnd`. This function cannot be used with `threshacceptbnd`.
- `@acceptancethresh` — Threshold acceptance acceptance function. The default for `threshacceptbnd`. This function cannot be used with `simulannealbnd`.
- `@myfun` — A custom acceptance function, `myfun`. See the functions above for syntax.

## Hybrid Function Options

A hybrid function is another minimization function that runs during or at the end of iterations of the solver. You can specify a hybrid function using the `HybridFcn` option. The choices are

- `[]` — No hybrid function
- `@fminsearch` — Uses the MATLAB function `fminsearch` to perform unconstrained minimization

- `@patternsearch` — Uses `patternsearch` to perform constrained or unconstrained minimization
- `@fminunc` — Uses the Optimization Toolbox function `fminunc` to perform unconstrained minimization
- `@fmincon` — Uses the Optimization Toolbox function `fmincon` to perform constrained minimization

`HybridInterval` specifies the interval (if not `never` or `end`) at which the hybrid function is called.

## Stopping Criteria Options

Stopping criteria determine what causes the algorithm to terminate. You can specify the following options:

- `TolFun` — The algorithm runs until the average change in value of the objective function in `StallIterLim` iterations is less than `TolFun`. The default value is `1e-6`.
- `MaxIter` — The algorithm stops if the number of iterations exceeds this maximum number of iterations. You can specify the maximum number of iterations as a positive integer or `Inf`. `Inf` is the default.
- `MaxFunEval` specifies the maximum number of evaluations of the objective function. The algorithm stops if the number of function evaluations exceeds the maximum number of function evaluations. The allowed maximum is `3000*numberofvariables`.
- `TimeLimit` specifies the maximum time in seconds the algorithm runs before stopping.
- `ObjectiveLimit` — The algorithm stops if the best objective function value is less than or equal to the value of `ObjectiveLimit`.

## Output Function Options

Output functions are functions that the algorithm calls at each iteration. The default value is to have no output function, `[]`. You must first create an output function using the syntax described in “Structure of the Output Function” on page 9-51. Then, specify your function as `@myfun`, where `myfun` is the name of your function.



If you are using `simulanneal`, set

```
options = saoptimset('OutputFcns', @myfun);
```

To see a template that you can use to write your own output functions, enter

```
edit saoutputfcntemplate
```

at the MATLAB command line.

### Structure of the Output Function

The output function has the following calling syntax.

```
[stop, options, optchanged] = myfun(options, optimvalues, flag)
```

The function has the following input arguments:

- `options` — Options structure created using `saoptimset`.
- `optimvalues` — Structure containing information about the current state of the solver. The structure contains the following fields:
  - `x` — Current point
  - `fval` — Objective function value at `x`
  - `bestx` — Best point found so far
  - `bestfval` — Objective function value at best point
  - `temperature` — Current temperature
  - `iteration` — Current iteration
  - `funccount` — Number of function evaluations
  - `t0` — Start time for algorithm
  - `k` — Annealing parameter
  - `tau` — Threshold acceptance sequence (for threshold acceptance solver only)
- `flag` — Current state in which the output function is called. The possible values for `flag` are

- 'init' — Initialization state
- 'iter' — Iteration state
- 'done' — Final state

The output function returns the following arguments:

- stop — Provides a way to stop the algorithm at the current iteration. stop can have the following values:
  - false — The algorithm continues to the next iteration.
  - true — The algorithm terminates at the current iteration.
- options — Options structure modified by the output function.
- optchanged — A boolean flag indicating changes were made to options. This must be set to true if options are changed.

## Display Options

Use the Display option to specify how much information is displayed at the command line while the algorithm is running. The available options are

- off — No output is displayed. This is the default value for an options structure created using `saoptimset`.
- iter — Information is displayed at each iteration.
- diagnose — Information is displayed at each iteration. In addition, the diagnostic lists some problem information and the options that have been changed from the defaults.
- final — The reason for stopping is displayed. This is the default.

Both iter and diagnose display the following information:

- Iteration — Iteration number
- f-count — Cumulative number of objective function evaluations
- Best  $f(x)$  — Best objective function value
- Current  $f(x)$  — Current objective function value
- Mean Temperature — Mean temperature function value

# Functions — By Category

---

Genetic Algorithm (p. 10-2)

Use genetic algorithm and Genetic Algorithm Tool, and modify genetic algorithm options

Direct Search (p. 10-2)

Use direct search and Pattern Search Tool, and modify pattern search options

Simulated Annealing and Threshold Acceptance Algorithms (p. 10-2)

Use simulated annealing and threshold acceptance algorithms, and modify algorithm options

## Genetic Algorithm

|                         |  |
|-------------------------|--|
| <code>ga</code>         | Find minimum of function using genetic algorithm |
| <code>gaoptimget</code> | Values of genetic algorithm options structure    |
| <code>gaoptimset</code> | Create genetic algorithm options structure       |
| <code>gatool</code>     | Open Genetic Algorithm Tool                      |

## Direct Search

|                            |   |
|----------------------------|---|
| <code>patternsearch</code> | Find minimum of function using pattern search |
| <code>psearchtool</code>   | Open Pattern Search Tool                      |
| <code>psoptimget</code>    | Values of pattern search options structure    |
| <code>psoptimset</code>    | Create pattern search options structure       |

## Simulated Annealing and Threshold Acceptance Algorithms

|                         |   |
|-------------------------|---|
| <code>saoptimget</code> | Values of simulated annealing or threshold acceptance algorithm options structure |
| <code>saoptimset</code> | Create simulated annealing algorithm or threshold acceptance options structure    |

`simulannealbnd`

Find unconstrained or bound-constrained minimum of function of several variables using simulated annealing algorithm

`threshacceptbnd`

Find unconstrained or bound-constrained minimum of function of several variables using threshold acceptance algorithm



# Functions — Alphabetical List

---

**Purpose**

Find minimum of function using genetic algorithm

**Syntax**

```
x = ga(fitnessfcn,nvars)
x = ga(fitnessfcn,nvars,A,b)
x = ga(fitnessfcn,nvars,A,b,Aeq,beq)
x = ga(fitnessfcn,nvars,A,b,Aeq,beq,LB,UB)
x = ga(fitnessfcn,nvars,A,b,Aeq,beq,LB,UB,nonlcon)
x = ga(fitnessfcn,nvars,A,b,Aeq,beq,LB,UB,nonlcon,options)
x = ga(problem)
[x, fval] = ga(...)
[x, fval, exitflag] = ga(...)
```

**Description**

ga implements the genetic algorithm at the command line to minimize an objective function.

`x = ga(fitnessfcn,nvars)` finds a local unconstrained minimum, `x`, to the objective function, `fitnessfcn`. `nvars` is the dimension (number of design variables) of `fitnessfcn`. The objective function, `fitnessfcn`, accepts a vector `x` of size 1-by-`nvars`, and returns a scalar evaluated at `x`.

---

**Note** To write a function with additional parameters to the independent variables that can be called by `ga`, see the section on “Parameterizing Functions Called by `ga`” on page 7-67.

---

`x = ga(fitnessfcn,nvars,A,b)` finds a local minimum `x` to `fitnessfcn`, subject to the linear inequalities  $A * x \leq b$ . `fitnessfcn` accepts input `x` and returns a scalar function value evaluated at `x`.

If the problem has `m` linear inequality constraints and `n` variables, then

- `A` is a matrix of size `m`-by-`n`.
- `b` is a vector of length `m`.

Note that the linear constraints are not satisfied when the `PopulationType` option is set to `'bitString'` or `'custom'`.



`x = ga(fitnessfcn,nvars,A,b,Aeq,beq)` finds a local minimum  $x$  to `fitnessfcn`, subject to the linear equalities  $Aeq * x = beq$  as well as  $A * x \leq b$ . (Set `A=[]` and `b=[]` if no inequalities exist.)

If the problem has  $r$  linear equality constraints and  $n$  variables, then

- `Aeq` is a matrix of size  $r$ -by- $n$ .
- `beq` is a vector of length  $r$ .

Note that the linear constraints are not satisfied when the `PopulationType` option is set to `'bitString'` or `'custom'`.

`x = ga(fitnessfcn,nvars,A,b,Aeq,beq,LB,UB)` defines a set of lower and upper bounds on the design variables,  $x$ , so that a solution is found in the range  $LB \leq x \leq UB$ . Use empty matrices for `LB` and `UB` if no bounds exist. Set  $LB(i) = -Inf$  if  $x(i)$  is unbounded below; set  $UB(i) = Inf$  if  $x(i)$  is unbounded above.

`x = ga(fitnessfcn,nvars,A,b,Aeq,beq,LB,UB,nonlcon)` subjects the minimization to the constraints defined in `nonlcon`. The function `nonlcon` accepts  $x$  and returns the vectors `C` and `Ceq`, representing the nonlinear inequalities and equalities respectively. `ga` minimizes the `fitnessfcn` such that  $C(x) \leq 0$  and  $Ceq(x) = 0$ . (Set `LB=[]` and `UB=[]` if no bounds exist.)

Note that the nonlinear constraints are not satisfied when the `PopulationType` option is set to `'bitString'` or `'custom'`.

`x = ga(fitnessfcn,nvars,A,b,Aeq,beq,LB,UB,nonlcon,options)` minimizes with the default optimization parameters replaced by values in the structure `options`, which can be created using the `gaoptimset` function. See the `gaoptimset` reference page for details.

`x = ga(problem)` finds the minimum for `problem`, where `problem` is a structure containing the following fields:

|                         |                            |
|-------------------------|----------------------------|
| <code>fitnessfcn</code> | Fitness function           |
| <code>nvars</code>      | Number of design variables |

|            |   |
|------------|---|
| options    | Options structure created using <code>gaoptimset</code> |
| Aineq      | A matrix for inequality constraints                     |
| Bineq      | b vector for inequality constraints                     |
| Aeq        | A matrix for equality constraints                       |
| Beq        | b vector for equality constraints                       |
| LB         | Lower bound on x  |
| UB         | Upper bound on x  |
| nonlcon    | Nonlinear constraint function                           |
| randstate  | Optional field to reset rand state                      |
| randnstate | Optional field to reset randn state                     |

`[x, fval] = ga(...)` returns `fval`, the value of the fitness function at `x`.

`[x, fval, exitflag] = ga(...)` returns `exitflag`, an integer identifying the reason the algorithm terminated. The following lists the values of `exitflag` and the corresponding reasons the algorithm terminated.

- 1 — Average cumulative change in value of the fitness function over `options.StallGenLimit` generations less than `options.TolFun` and constraint violation less than `options.TolCon`.
- 2 — Fitness limit reached and constraint violation less than `options.TolCon`.
- 3 — The value of the fitness function did not change in `options.StallGenLimit` generations and constraint violation less than `options.TolCon`.
- 4 — Magnitude of step smaller than machine precision and constraint violation less than `options.TolCon`.
- 0 — Maximum number of generations exceeded.
- -1 — Optimization terminated by the output or plot function.

- -2 — No feasible point found.
- -4 — Stall time limit exceeded.
- -5 — Time limit exceeded.

`[x, fval, exitflag, output] = ga(...)` returns `output`, a structure that contains output from each generation and other information about the performance of the algorithm. The output structure contains the following fields:

- `randstate` — The state of `rand`, the MATLAB random number generator, just before the algorithm started.
- `randnstate` — The state of `randn` the MATLAB normal random number generator, just before the algorithm started. You can use the values of `randstate` and `randnstate` to reproduce the output of `ga`. See “Reproducing Your Results” on page 7-29.
- `generations` — The number of generations computed.
- `funccount` — The number of evaluations of the fitness function
- `message` — The reason the algorithm terminated.
- `maxconstraint` — Maximum constraint violation, if any.

`[x, fval, exitflag, output, population] = ga(...)` returns the matrix, `population`, whose rows are the final population.

`[x, fval, exitflag, output, population, scores] = ga(...)` returns `scores` the scores of the final population.

---

**Note** For problems that use the population type `Double Vector` (the default), `ga` does not accept functions whose inputs are of type `complex`. To solve problems involving complex data, write your functions so that they accept real vectors, by separating the real and imaginary parts.

---

**Example**

Given the following inequality constraints and lower bounds

$$\begin{bmatrix} 1 & 1 \\ -1 & 2 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \leq \begin{bmatrix} 2 \\ 2 \\ 3 \end{bmatrix}$$

$$\begin{bmatrix} 0 \leq x_1 \\ 0 \leq x_2 \end{bmatrix}$$

the following code finds the minimum of the function, `lincontest6`, that is provided with the toolbox:

```
A = [1 1; -1 2; 2 1];  
b = [2; 2; 3];  
lb = zeros(2,1);  
[x, fval, exitflag] = ga(@lincontest6,...  
2,A,b,[],[],lb)
```

Optimization terminated:

average change in the fitness value less than options.TolFun.

x =

```
0.7794  1.2205
```

fval =

```
-8.03916
```

exitflag =

```
1
```

**References**

[1] Goldberg, David E., *Genetic Algorithms in Search, Optimization & Machine Learning*, Addison-Wesley, 1989.

[2] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Algorithm for Optimization with General Constraints and Simple Bounds”, *SIAM Journal on Numerical Analysis*, Volume 28, Number 2, pages 545–572, 1991.

[3] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Barrier Algorithm for Optimization with General Inequality Constraints and Simple Bounds”, *Mathematics of Computation*, Volume 66, Number 217, pages 261–288, 1997.

**See Also**

goptimset, gatool, patternsearch, simulannealbnd,  
threshacceptbnd

# gaoptimget

---

**Purpose** Values of genetic algorithm options structure

**Syntax**  
`val = gaoptimget(options, 'name')`  
`val = gaoptimget(options, 'name', default)`

**Description** `val = gaoptimget(options, 'name')` returns the value of the parameter name from the genetic algorithm options structure options. `gaoptimget(options, 'name')` returns an empty matrix [] if the value of name is not specified in options. It is only necessary to type enough leading characters of name to uniquely identify it. `gaoptimget` ignores case in parameter names.

`val = gaoptimget(options, 'name', default)` returns the 'name' parameter, but will return the default value if the name parameter is not specified (or is []) in options.

**See Also** `ga`, `gaoptimset`, `gatool`

**Purpose** Create genetic algorithm options structure

**Syntax**

```
gaoptimset
options = gaoptimset
options = gaoptimset('param1',value1,'param2',value2,...)
options = gaoptimset(olddopts,'param1',value1,...)
options = gaoptimset(olddopts,newopts)
```

**Description** gaoptimset with no input or output arguments displays a complete list of parameters with their valid values.

options = gaoptimset (with no input arguments) creates a structure called options that contains the options, or *parameters*, for the genetic algorithm and sets parameters to their default values.

options = gaoptimset('param1',value1,'param2',value2,...) creates a structure options and sets the value of 'param1' to value1, 'param2' to value2, and so on. Any unspecified parameters are set to their default values. It is sufficient to type only enough leading characters to define the parameter name uniquely. Case is ignored for parameter names.

options = gaoptimset(olddopts,'param1',value1,...) creates a copy of olddopts, modifying the specified parameters with the specified values.

options = gaoptimset(olddopts,newopts) combines an existing options structure, olddopts, with a new options structure, newopts. Any parameters in newopts with nonempty values overwrite the corresponding old parameters in olddopts.

**Options** The following table lists the options you can set with gaoptimset. See “Importing and Exporting Options and Problems” on page 6-16 for a complete description of these options and their values. Values in {} denote the default value. You can also view the optimization parameters and defaults by typing gaoptimset at the command line.

# gaoptimset

| Option            | Description  | Values  |
|-------------------|--|---|
| CreationFcn       | Handle to the function that creates the initial population   | {@gacreationuniform}  |
| CrossoverFcn      | Handle to the function that the algorithm uses to create crossover children  | @crossoverheuristic<br>{@crossoverscattered}<br>@crossoverintermediate<br>@crossoversinglepoint<br>@crossovertwopoint<br>@crossoverarithmetic |
| CrossoverFraction | The fraction of the population at the next generation, not including elite children, that is created by the crossover function | Positive scalar   {0.8}   |
| Display           | Level of display   | 'off'   'iter'   'diagnose'<br>  {'final'}  |
| EliteCount        | Positive integer specifying how many individuals in the current generation are guaranteed to survive to the next generation    | Positive integer   {2}  |
| FitnessLimit      | Scalar. If the fitness function attains the value of FitnessLimit, the algorithm halts.  | Scalar   {-Inf}   |
| FitnessScalingFcn | Handle to the function that scales the values of the fitness function  | @fitscalingshiftlinear<br>@fitscalingprop<br>@fitscalingtop<br>{@fitscalingrank}  |



| Option             | Description  | Values   |
|--------------------|--|--|
| Generations        | Positive integer specifying the maximum number of iterations before the algorithm halts  | Positive integer   {100}   |
| HybridFcn          | Handle to a function that continues the optimization after ga terminates   | Function handle<br>  @fminsearch<br>@patternsearch<br>@fminunc<br>@fmincon<br>{[]} |
| InitialPenalty     | Initial value of penalty parameter   | Positive scalar   {10}   |
| InitialPopulation  | Initial population used to seed the genetic algorithm  | Matrix   {{{}}   |
| InitialScores      | Initial scores used to determine fitness   | Column vector   {{{}}  |
| MigrationDirection | Direction of migration   | 'both'   {'forward'}   |
| MigrationFraction  | Scalar between 0 and 1 specifying the fraction of individuals in each subpopulation that migrates to a different subpopulation | Scalar   {0.2}   |
| MigrationInterval  | Positive integer specifying the number of generations that take place between migrations of individuals between subpopulations | Positive integer   {20}  |
| MutationFcn        | Handle to the function that produces mutation children   | @mutationuniform<br>@mutationadaptfeasible<br>{@mutationgaussian}                  |

# gaoptimset

| Option         | Description   | Values  |
|----------------|---|---|
| OutputFcns     | Functions that ga calls at each iteration   | @gaoutputgen   {}   |
| PenaltyFactor  | Penalty update parameter  | Positive scalar   {100}   |
| PlotFcns       | Array of handles to functions that plot data computed by the algorithm                                | @gaplotbestf<br>@gaplotbestindiv<br>@gaplotdistance<br>@gaplotexpectation<br>@gaplotgeneology<br>@gaplotselection<br>@gaplotrange<br>@gaplotcorediversity<br>@gaplotcores<br>@gaplotstopping   {} |
| PlotInterval   | Positive integer specifying the number of generations between consecutive calls to the plot functions | Positive integer   {1}  |
| PopInitRange   | Matrix or vector specifying the range of the individuals in the initial population                    | Matrix or vector   [0;1]  |
| PopulationSize | Size of the population  | Positive integer   {20}   |
| PopulationType | String describing the data type of the population   | 'bitstring'   'custom'   {'doubleVector'}<br><br>Note that linear and nonlinear constraints are not satisfied when PopulationType is set to 'bitString' or 'custom'.                              |

| Option         | Description   | Values  |
|----------------|---|---|
| SelectionFcn   | Handle to the function that selects parents of crossover and mutation children  | @selectionremainder<br>@selectionuniform<br>{@selectionstochunif}<br>@selectionroulette<br>@selectiontournament |
| StallGenLimit  | Positive integer. The algorithm stops if there is no improvement in the objective function for StallGenLimit consecutive generations. | Positive integer   {50}   |
| StallTimeLimit | Positive scalar. The algorithm stops if there is no improvement in the objective function for StallTimeLimit seconds.                 | Positive scalar   {20}  |
| TimeLimit      | Positive scalar. The algorithm stops after running for TimeLimit seconds.   | Positive scalar   {Inf}   |
| TolCon         | Positive scalar. TolCon is used to determine the feasibility with respect to nonlinear constraints.                                   | Positive scalar   {1e-6}  |

# gaoptimset

---

| Option     | Description   | Values                   |
|------------|---|--------------------------|
| TolFun     | Positive scalar. The algorithm runs until the cumulative change in the fitness function value over StallGenLimit is less than TolFun. | Positive scalar   {1e-6} |
| Vectorized | String specifying whether the computation of the fitness function is vectorized   | 'on'   {'off'}           |

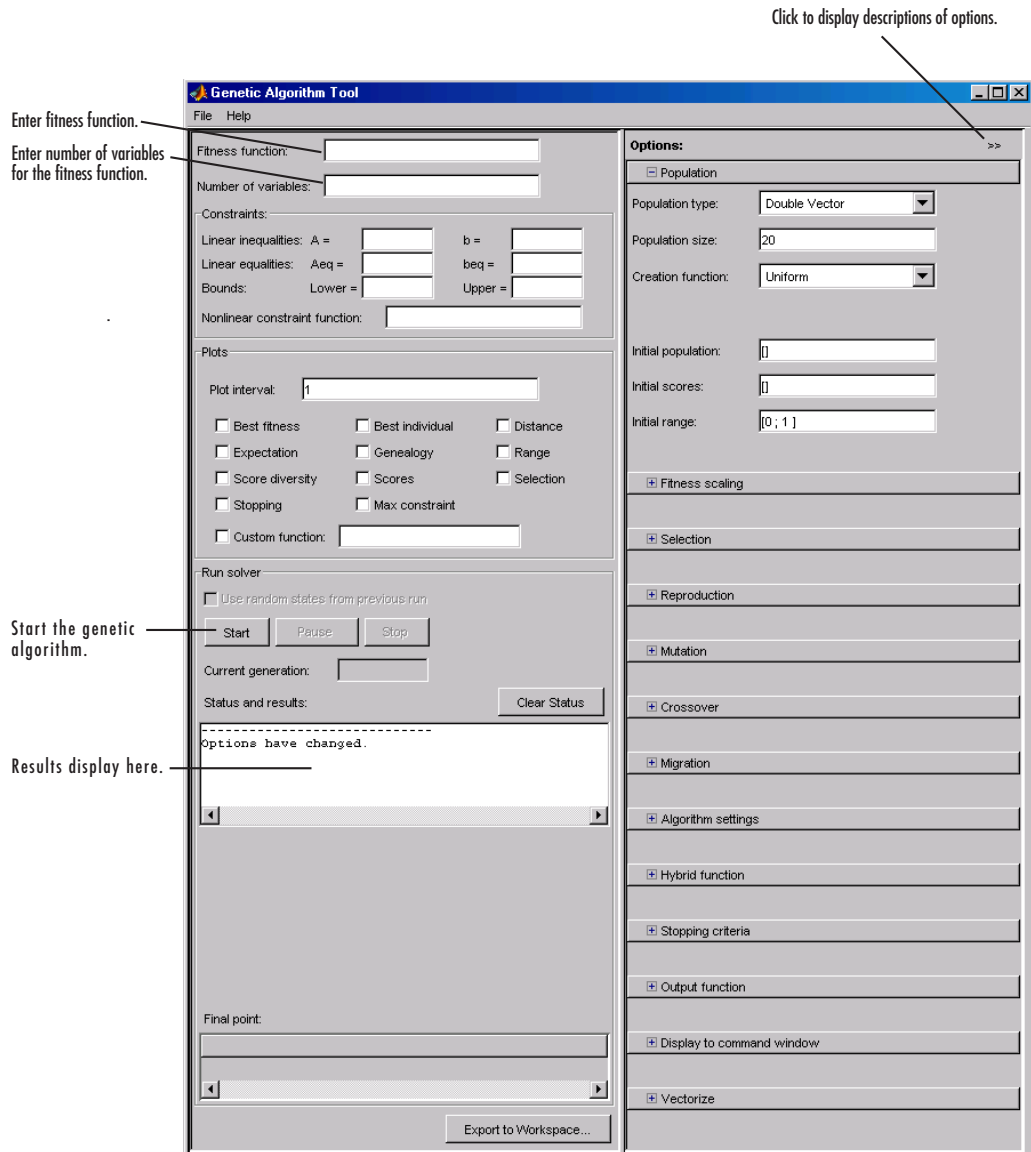
**See Also**      gaoptimget, gatool

**Purpose** Open Genetic Algorithm Tool

**Syntax** `gatool`  
`gatool(optstruct)`

## Description

gatool opens the Genetic Algorithm Tool, a graphical user interface (GUI) to the genetic algorithm, as shown in the figure below.



`gatool(optstruct)` starts the Genetic Algorithm Tool with `optstruct`, which can be either a genetic algorithm options structure or genetic algorithm problem structure. An options structure can be created using `gaoptimset`.

`gatool` can run the genetic algorithm solver, `ga` on optimization problems, and it allows you to modify several `ga` options at run time, while you watch your results. See “Using the Genetic Algorithm Tool” on page 4-4 for a complete description of the tool.

**See Also**

`ga`, `gaoptimset`, `psearchtool`

# patternsearch

---

**Purpose** Find minimum of function using pattern search

**Syntax**

```
x = patternsearch(@fun,x0)
x = patternsearch(@fun,x0,A,b)
x = patternsearch(@fun,x0,A,b,Aeq,beq)
x = patternsearch(@fun,x0,A,b,Aeq,beq,LB,UB)
x = patternsearch(@fun,x0,A,b,Aeq,beq,LB,UB,nonlcon)
x = patternsearch(@fun,x0,A,b,Aeq,beq,LB,UB,nonlcon,options)
x = patternsearch(problem)
[x, fval] = patternsearch(@fun,x0, ...)
[x, fval, exitflag] = patternsearch(@fun,x0, ...)
[x, fval, exitflag, output] = patternsearch(@fun,x0, ...)
```

**Description** patternsearch finds the minimum of a function using a pattern search. `x = patternsearch(@fun,x0)` finds the local minimum, `x`, to the MATLAB function, `fun`, that computes the values of the objective function  $f(x)$ , and `x0` is an initial point for the pattern search algorithm. The function `patternsearch` accepts the objective function as a function handle of the form `@fun`. The function `fun` accepts a vector input and returns a scalar function value.

---

**Note** To write a function with additional parameters to the independent variables that can be called by `patternsearch`, see the section on “Parameterizing Functions Called by `patternsearch`” on page 6-53.

---

`x = patternsearch(@fun,x0,A,b)` finds a local minimum `x` to the function `fun`, subject to the linear inequality constraints represented in matrix form by  $Ax \leq b$ .

If the problem has `m` linear inequality constraints and `n` variables, then

- `A` is a matrix of size `m-by-n`.
- `b` is a vector of length `m`.



`x = patternsearch(@fun,x0,A,b,Aeq,beq)` finds a local minimum  $x$  to the function `fun`, starting at `x0`, and subject to the constraints

$$Ax \leq b$$

$$Aeq * x = beq$$

where  $Aeq * x = beq$  represents the linear equality constraints in matrix form. If the problem has  $r$  linear equality constraints and  $n$  variables, then

- `Aeq` is a matrix of size  $r$ -by- $n$ .
- `beq` is a vector of length  $r$ .

If there are no inequality constraints, pass empty matrices, `[]`, for `A` and `b`.

`x = patternsearch(@fun,x0,A,b,Aeq,beq,LB,UB)` defines a set of lower and upper bounds on the design variables,  $x$ , so that a solution is found in the range  $LB \leq x \leq UB$ . If the problem has  $n$  variables, `LB` and `UB` are vectors of length  $n$ . If `LB` or `UB` is empty (or not provided), it is automatically expanded to `-Inf` or `Inf`, respectively. If there are no inequality or equality constraints, pass empty matrices for `A`, `b`, `Aeq` and `beq`.

`x = patternsearch(@fun,x0,A,b,Aeq,beq,LB,UB,nonlcon)` subjects the minimization to the constraints defined in `nonlcon`, a nonlinear constraint function. The function `nonlcon` accepts  $x$  and returns the vectors `C` and `Ceq`, representing the nonlinear inequalities and equalities respectively. `fmincon` minimizes `fun` such that  $C(x) \leq 0$  and  $Ceq(x)=0$ . (Set `LB=[]` and `UB=[]` if no bounds exist.)

`x = patternsearch(@fun,x0,A,b,Aeq,beq,LB,UB,nonlcon,options)` minimizes `fun` with the default optimization parameters replaced by values in `options`. The structure `options` can be created using `psoptimset`.

# patternsearch

---

`x = patternsearch(problem)` finds the minimum for `problem`, where `problem` is a structure containing the following fields:

- `objective` — Objective function
- `X0` — Starting point
- `Aineq` — Matrix for the inequality constraints
- `Bineq` — Vector for the inequality constraints
- `Aeq` — Matrix for the equality constraints
- `Beq` — Vector for the equality constraints
- `LB` — Lower bound for `x`
- `UB` — Upper bound for `x`
- `nonlcon` — Nonlinear constraint function
- `options` — Options structure created with `psoptimset`
- `randstate` — Optional field to reset the state of `rand`
- `randnstate` — Optional field to reset the state of `randn`

You can create the structure `problem` by exporting a problem from the Pattern Search Tool, as described in “Importing and Exporting Options and Problems” on page 6-16.

---

**Note** `problem` must have all the fields as specified above.

---

`[x, fval] = patternsearch(@fun,x0, ...)` returns the value of the objective function `fun` at the solution `x`.

`[x, fval, exitflag] = patternsearch(@fun,x0, ...)` returns `exitflag`, which describes the exit condition of `patternsearch`. Possible values of `exitflag` and the corresponding conditions are

- |    |   |
|----|---|
| 1  | Magnitude of mesh size is less than specified tolerance and constraint violation less than <code>options.TolCon</code> .    |
| 2  | Change in <code>x</code> less than the specified tolerance and constraint violation less than <code>options.TolCon</code> . |
| 4  | Magnitude of step smaller than machine precision and constraint violation less than <code>options.TolCon</code> .           |
| 0  | Maximum number of function evaluations or iterations reached.   |
| -1 | Optimization terminated by the output or plot function.   |
| -2 | No feasible point found.  |

`[x, fval, exitflag, output] = patternsearch(@fun,x0, ...)`  
returns a structure output containing information about the search.  
The output structure contains the following fields:

- `function` — Objective function
- `problemtype` — Type of problem: unconstrained, bound constrained or linear constrained
- `pollmethod` — Polling technique
- `searchmethod` — Search technique used, if any
- `iteration` — Total number of iterations
- `funccount` — Total number of function evaluations
- `meshsize` — Mesh size at `x`
- `maxconstraint` — Maximum constraint violation, if any
- `message` — Reason why the algorithm terminated

---

**Note** patternsearch does not accept functions whose inputs are of type complex. To solve problems involving complex data, write your functions so that they accept real vectors, by separating the real and imaginary parts.

---

## Example

Given the following constraints

$$\begin{bmatrix} 1 & 1 \\ -1 & 2 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \leq \begin{bmatrix} 2 \\ 2 \\ 3 \end{bmatrix}$$

$$\begin{bmatrix} 0 \leq x_1 \\ 0 \leq x_2 \end{bmatrix}$$

the following code finds the minimum of the function, lincontest6, that is provided with the toolbox:

```
A = [1 1; -1 2; 2 1];
b = [2; 2; 3];
lb = zeros(2,1);
[x, fval, exitflag] = patternsearch(@lincontest6,[0 0],...
                                   A,b,[],[],lb)
Optimization terminated: mesh size less than
                        options.TolMesh.
```

```
x =
```

```
    0.6667    1.3333
```

```
fval =
```

```
   -8.2222
```

exitflag =

1

## References

- [1] Torczon, Virginia, “On the Convergence of Pattern Search Algorithms”, *SIAM Journal on Optimization*, Volume 7, Number 1, pages 1–25, 1997.
- [2] Lewis, Robert Michael and Virginia Torczon, “Pattern Search Algorithms for Bound Constrained Minimization”, *SIAM Journal on Optimization*, Volume 9, Number 4, pages 1082–1099, 1999.
- [3] Lewis, Robert Michael and Virginia Torczon, “Pattern Search Methods for Linearly Constrained Minimization”, *SIAM Journal on Optimization*, Volume 10, Number 3, pages 917–941, 2000.
- [4] Audet, Charles and J. E. Dennis Jr., “Analysis of Generalized Pattern Searches”, *SIAM Journal on Optimization*, Volume 13, Number 3, pages 889–903, 2003.
- [5] Lewis, Robert Michael and Virginia Torczon, “A Globally Convergent Augmented Lagrangian Pattern Search Algorithm for Optimization with General Constraints and Simple Bounds”, *SIAM Journal on Optimization*, Volume 12, Number 4, pages 1075–1089, 2002.
- [6] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Pattern Search Algorithm for Optimization with General Constraints and Simple Bounds”, *SIAM Journal on Numerical Analysis*, Volume 28, Number 2, pages 545–572, 1991.
- [7] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. “A Globally Convergent Augmented Lagrangian Pattern Search Algorithm for Optimization with General Constraints and Simple Bounds”, *Mathematics of Computation*, Volume 66, Number 217, pages 261–288, 1997.

# patternsearch

---

## **See Also**

psearchtool, psoptimget, psoptimset, ga, simulannealbnd,  
threshacceptbnd

**Purpose** Open Pattern Search Tool

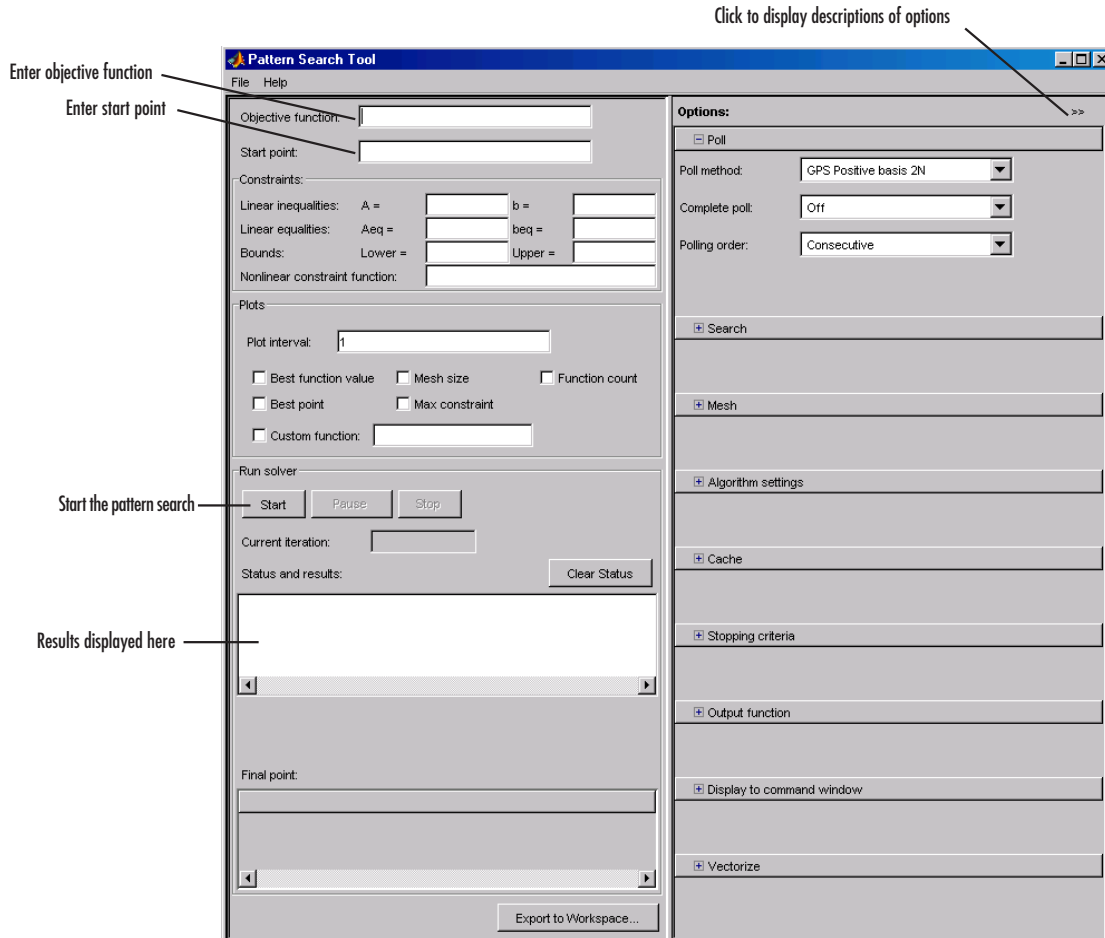
**Syntax** `psearchtool`  
`psearchtool(optstruct)`

**Description** `psearchtool` opens the Pattern Search Tool, a graphical user interface (GUI) for performing pattern searches, as shown in the figure below. You can use the Pattern Search Tool to run a pattern search on optimization problems and display the results. See “Using the Pattern Search Tool” on page 3-3 for a complete description of the tool.

`psearchtool(optstruct)` starts the Pattern Search tool using `optstruct`, which can be a Pattern Search options structure or Pattern Search problem structure. An options structure can be created using `psoptimset`.

# psearchtool

psearchtool can run the Pattern Search Algorithm solver, patternsearch on optimization problems, and it allows you to modify several patternsearch options at run time, while you watch your results.



## See Also

patternsearch, psoptimget, psoptimset, gatool



**Purpose** Values of pattern search options structure

**Syntax**  
`val = psoptimget(options, 'name')`  
`val = psoptimget(options, 'name', default)`

**Description**  
`val = psoptimget(options, 'name')` returns the value of the parameter name from the pattern search options structure options. `psoptimget(options, 'name')` returns an empty matrix [] if the value of name is not specified in options. It is only necessary to type enough leading characters of name to uniquely identify it. `psoptimget` ignores case in parameter names.  
`val = psoptimget(options, 'name', default)` returns the value of the parameter name from the pattern search options structure options, but returns default if the parameter is not specified (as in []) in options.

**Example**  
`val = psoptimget(opts, 'TolX', 1e-4);`

returns `val = 1e-4` if the TolX property is not specified in opts.

**See Also** `psoptimset`, `patternsearch`

# psoptimset

---

**Purpose** Create pattern search options structure

**Syntax**

```
psoptimset
options = psoptimset
options = psoptimset('param1',value1,'param2',value2,...)
options = psoptimset(oldopts,'param1',value1,...)
options = psoptimset(oldopts,newopts)
```

**Description** psoptimset with no input or output arguments displays a complete list of parameters with their valid values.

options = psoptimset (with no input arguments) creates a structure called options that contains the options, or *parameters*, for the pattern search and sets parameters to their default values.

options = psoptimset('param1',value1,'param2',value2,...) creates a structure options and sets the value of 'param1' to value1, 'param2' to value2, and so on. Any unspecified parameters are set to their default values. It is sufficient to type only enough leading characters to define the parameter name uniquely. Case is ignored for parameter names.

options = psoptimset(oldopts,'param1',value1,...) creates a copy of oldopts, modifying the specified parameters with the specified values.

options = psoptimset(oldopts,newopts) combines an existing options structure, oldopts, with a new options structure, newopts. Any parameters in newopts with nonempty values overwrite the corresponding old parameters in oldopts.

**Options** The following table lists the options you can set with psoptimset. See for a complete description of the options and their values. Values in {} denote the default value. You can also view the optimization parameters and defaults by typing psoptimset at the command line.

| Option          | Description  | Values                                  |
|-----------------|--|---|
| Cache           | With Cache set to 'on', patternsearch keeps a history of the mesh points it polls and does not poll points close to them again at subsequent iterations. Use this option if patternsearch runs slowly because it is taking a long time to compute the objective function. If the objective function is stochastic, it is advised not to use this option. | 'on'   {'off'}                          |
| CacheSize       | Size of the history  | Positive scalar   {1e4}                 |
| CacheTol        | Positive scalar specifying how close the current mesh point must be to a point in the history in order for patternsearch to avoid polling it. Use if 'Cache' option is set to 'on'.  | Positive scalar   {eps}                 |
| CompletePoll    | Complete poll around current iterate   | 'on'   {'off'}                          |
| CompleteSearch  | Complete poll around current iterate   | 'on'   {'off'}                          |
| Display         | Level of display   | 'off'   'iter'   'diagnose'   {'final'} |
| InitialMeshSize | Initial mesh size for pattern algorithm  | Positive scalar   {1.0}                 |
| InitialPenalty  | Initial value of the penalty parameter   | Positive scalar   {10}                  |

# psoptimset

| Option          | Description   | Values   |
|-----------------|---|--|
| MaxFunEvals     | Maximum number of objective function evaluations  | Positive integer   {2000*numberOfVariables}                          |
| MaxIter         | Maximum number of iterations  | Positive integer   {100*numberOfVariables}                           |
| MaxMeshSize     | Maximum mesh size used in a poll/search step  | Positive scalar   {Inf}  |
| MeshAccelerator | Accelerate convergence near a minimum   | 'on'   {'off'}   |
| MeshContraction | Mesh contraction factor, used when iteration is unsuccessful                            | Positive scalar   {0.5}  |
| MeshExpansion   | Mesh expansion factor, expands mesh when iteration is successful                        | Positive scalar   {2.0}  |
| MeshRotate      | Rotate the pattern before declaring a point to be optimum                               | 'on'   {'off'}   |
| OutputFcn       | Specifies a user-defined function that an optimization function calls at each iteration | @psoutputhistory   {}  |
| PenaltyFactor   | Penalty update parameter  | Positive scalar   {100}  |
| PlotFcn         | Specifies plots of output from the pattern search                                       | @psplotbestf   @psplotmeshsize   @psplotfuncount   @psplotbestx   {} |
| PlotInterval    | Specifies that plot functions will be called at every interval                          | {1}  |

| Option       | Description                                      | Values  |
|--------------|--|---|
| PollingOrder | Order of poll directions in pattern search       | 'Random'   'Success'   {'Consecutive'}  |
| PollMethod   | Polling strategy used in pattern search          | {'GPSPositiveBasis2N'}   'GPSPositiveBasisNp1'   'MADSPositiveBasis2N'   'MADSPositiveBasisNp1'   |
| ScaleMesh    | Automatic scaling of variables                   | {'on'}   'off'  |
| SearchMethod | Type of search used in pattern search            | 'GPSPositiveBasisNp1'   'GPSPositiveBasis2N'   'MADSPositiveBasisNp1'   'MADSPositiveBasis2N'   @searchga   @searchlhs   @searchneldermead   {} |
| TimeLimit    | Total time (in seconds) allowed for optimization | Positive scalar   {Inf}   |
| TolBind      | Binding tolerance                                | Positive scalar   {1e-3}  |
| TolCon       | Tolerance on constraints                         | Positive scalar   {1e-6}  |
| TolFun       | Tolerance on function                            | Positive scalar   {1e-6}  |
| TolMesh      | Tolerance on mesh size                           | Positive scalar   {1e-6}  |
| TolX         | Tolerance on variable                            | Positive scalar   {1e-6}  |
| Vectorized   | Specifies whether functions are vectorized       | 'on'   {'off'}  |

For a detailed description of these options, see .

**See Also** patternsearch, psoptimget

# simulannealbnd

---

**Purpose** Find unconstrained or bound-constrained minimum of function of several variables using simulated annealing algorithm

**Syntax**

```
x = simulannealbnd(fun,x0)
x = simulannealbnd(fun,x0,lb,ub)
x = simulannealbnd(fun,x0,lb,ub,options)
[x,fval] = simulannealbnd(...)
[x,fval,exitflag] = simulannealbnd(...)
[x,fval,exitflag,output] = simulannealbnd(fun,...)
```

**Description** `x = simulannealbnd(fun,x0)` starts at `x0` and finds a local minimum `x` to the objective function specified by the function handle `fun`. The objective function accepts input `x` and returns a scalar function value evaluated at `x`. `x0` may be a scalar or a vector.

`x = simulannealbnd(fun,x0,lb,ub)` defines a set of lower and upper bounds on the design variables, `x`, so that a solution is found in the range  $lb \leq x \leq ub$ . Use empty matrices for `lb` and `ub` if no bounds exist. Set `lb(i)` to `-Inf` if `x(i)` is unbounded below; set `ub(i)` to `Inf` if `x(i)` is unbounded above.

`x = simulannealbnd(fun,x0,lb,ub,options)` minimizes with the default optimization parameters replaced by values in the structure `options`, which can be created using the `saoptimset` function. See the `saoptimset` reference page for details.

`[x,fval] = simulannealbnd(...)` returns `fval`, the value of the objective function at `x`.

`[x,fval,exitflag] = simulannealbnd(...)` returns `exitflag`, an integer identifying the reason the algorithm terminated. The following lists the values of `exitflag` and the corresponding reasons the algorithm terminated:

- 1 — Average change in the value of the objective function over `options.StallIterLimit` iterations is less than `options.TolFun`.
- 5 — `options.ObjectiveLimit` limit reached.
- 0 — Maximum number of function evaluations or iterations exceeded.

- -1 — Optimization terminated by the output or plot function.
- -2 — No feasible point found.
- -5 — Time limit exceeded.

`[x,fval,exitflag,output] = simulannealbnd(fun,...)` returns `output`, a structure that contains information about the problem and the performance of the algorithm. The output structure contains the following fields:

- `problemtype` — Type of problem: unconstrained or bound constrained.
- `iterations` — The number of iterations computed.
- `funcccount` — The number of evaluations of the objective function.
- `message` — The reason the algorithm terminated.
- `temperature` — Temperature when the solver terminated.
- `totaltime` — Total time for the solver to run.
- `randstate` — The state of `rand`, the MATLAB random number generator, just before the algorithm started.
- `randnstate` — The state of `randn` the MATLAB normal random number generator, just before the algorithm started. You can use the values of `randstate` and `randnstate` to reproduce the output of `simulannealbnd`. See “Reproducing Your Results” on page 8-5.


## Examples

Minimization of De Jong’s fifth function, a two-dimensional function with many local minima:

```
dejong5fcn
```

# simulannealbnd

---



```
x0 = [0 0];  
[x,fval] = simulannealbnd(@dejong5fcn,x0)
```

Optimization terminated: change in best function value  
less than options.TolFun.

```
x =  
    0.0392  -31.9700
```

```
fval =  
    2.9821
```

Minimization of De Jong's fifth function subject to lower and upper bounds:



```
x0 = [0 0];
lb = [-64 -64];
ub = [64 64];
[x,fval] = simulannealbnd(@dejong5fcn,x0,lb,ub)

Optimization terminated: change in best function value
                        less than options.TolFun.

x =
   -31.9652   -32.0286

fval =
    0.9980
```

The objective can also be an anonymous function:

```
fun = @(x) 3*sin(x(1))+exp(x(2));
x = simulannealbnd(fun,[1;1],[0 0])

Optimization terminated: change in best function value
                        less than options.TolFun.

x =
   457.1045
    0.0000
```

Minimization of De Jong's fifth function while displaying plots:

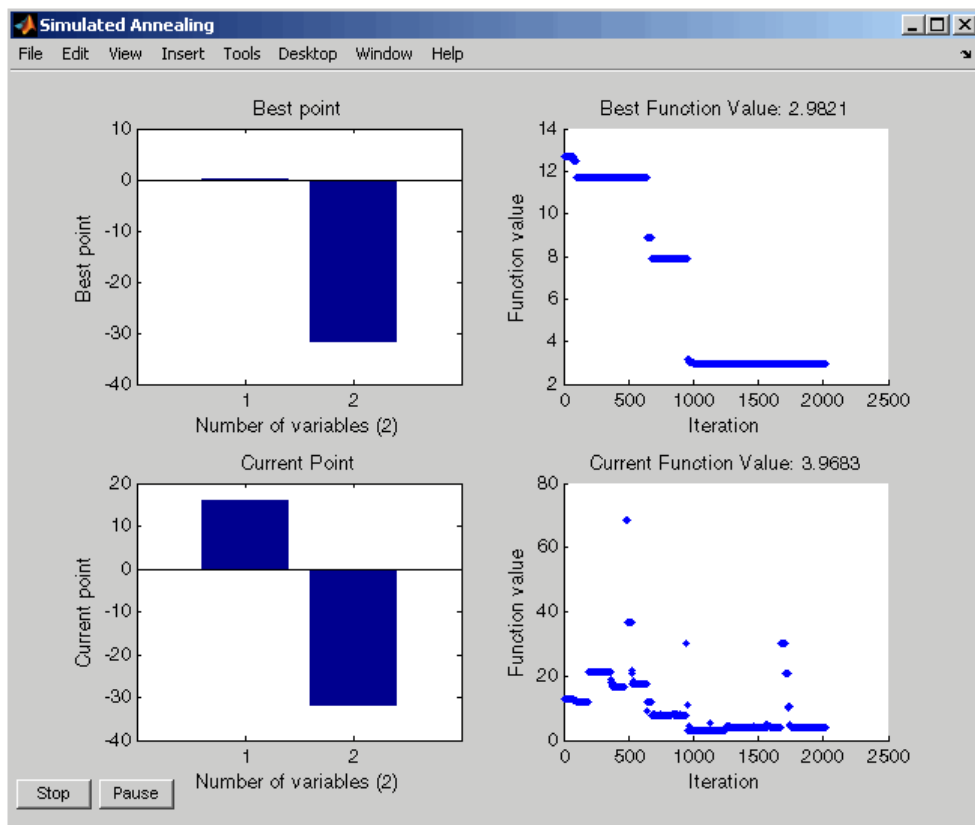
```
x0 = [0 0];
options = saoptimset('PlotFcns',{@saplotbestx,...
                            @saplotbestf,@saplotx,@saplotf});
simulannealbnd(@dejong5fcn,x0,[],[],options)

Optimization terminated: change in best function value
                        less than options.TolFun.
```

# simulannealbnd

```
ans =  
    0.0230 -31.9806
```

The plots displayed are shown below.



## See Also

saoptimset, threshacceptbnd, patternsearch, ga

## Purpose

Find unconstrained or bound-constrained minimum of function of several variables using threshold acceptance algorithm

## Syntax

```
x = threshacceptbnd(fun,x0)
x = threshacceptbnd(fun,x0,lb,ub)
x = threshacceptbnd(fun,x0,lb,ub,options)
[x,fval] = threshacceptbnd(...)
[x,fval,exitflag] = threshacceptbnd(...)
[x,fval,exitflag,output] = threshacceptbnd(...)
```

## Description

`x = threshacceptbnd(fun,x0)` starts at `x0` and finds a local minimum `x` to the objective function specified by the function handle `fun`. The objective function accepts input `x` and returns a scalar function value evaluated at `x`. `x0` may be a scalar or a vector.

`x = threshacceptbnd(fun,x0,lb,ub)` defines a set of lower and upper bounds on the design variables, `x`, so that a solution is found in the range  $lb \leq x \leq ub$ . Use empty matrices for `lb` and `ub` if no bounds exist. Set `lb(i)` to `-Inf` if `x(i)` is unbounded below; set `ub(i)` to `Inf` if `x(i)` is unbounded above.

`x = threshacceptbnd(fun,x0,lb,ub,options)` minimizes with the default optimization parameters replaced by values in the structure `options`, which can be created using the `soptimset` function. See the `soptimset` reference page for details.

`[x,fval] = threshacceptbnd(...)` returns `fval`, the value of the objective function at `x`.

`[x,fval,exitflag] = threshacceptbnd(...)` returns `exitflag`, an integer identifying the reason the algorithm terminated. The following lists the values of `exitflag` and the corresponding reasons the algorithm terminated:

- 1 — Average change in the value of the objective function over `options.StallIterLimit` iterations is less than `options.TolFun`.
- 5 — `options.ObjectiveLimit` limit reached.
- 0 — Maximum number of function evaluations or iterations exceeded.

# thresacceptbnd

---

- -1 — Optimization terminated by the output or plot function.
- -2 — No feasible point found.
- -5 — Time limit exceeded.


`[x,fval,exitflag,output] = thresacceptbnd(...)` returns `output`, a structure that contains information about the problem and the performance of the algorithm. The output structure contains the following fields:

- `problemtype` — Type of problem: unconstrained or bound constrained.
- `iterations` — The number of iterations computed.
- `funccount` — The number of evaluations of the objective function.
- `message` — The reason the algorithm terminated.
- `temperature` — Temperature when the solver terminated.
- `totaltime` — Total time for the solver to run.
- `randstate` — The state of `rand`, the MATLAB random number generator, just before the algorithm started.
- `randnstate` — The state of `randn`, the MATLAB normal random number generator, just before the algorithm started. You can use the values of `randstate` and `randnstate` to reproduce the output of `simulannealbnd`. See “Reproducing Your Results” on page 8-5.

## Examples

Minimizing De Jong’s fifth function, a two-dimensional function with many local minima:

```
dejong5fcn
```



```
x0 = [0 0];  
x = threshacceptbnd(@dejong5fcn,x0)
```

```
Optimization terminated: change in best function value  
less than options.TolFun.
```

```
x =  
    0    0
```

Minimization of De Jong's fifth function subject to lower and upper bounds:

```
x0 = [0 0];  
lb = [-64 -64]; ub = [64 64];
```

# threshacceptbnd

---

```
[x,fval] = threshacceptbnd(@dejong5fcn,x0,lb,ub)
```

```
Optimization terminated: change in best function value  
less than options.TolFun.
```

```
x =  
15.1309 -32.3303
```

```
fval =  
4.3932
```

Minimization of De Jong's fifth function while displaying plots:

```
x0 = [0 0];  
options = saoptimset('PlotFcns',{@saplotbestx,...  
@saplotbestf, @saplotx,@saplotf});  
[x,fval,exitflag] = threshacceptbnd(fun,x0,[],[],options)
```

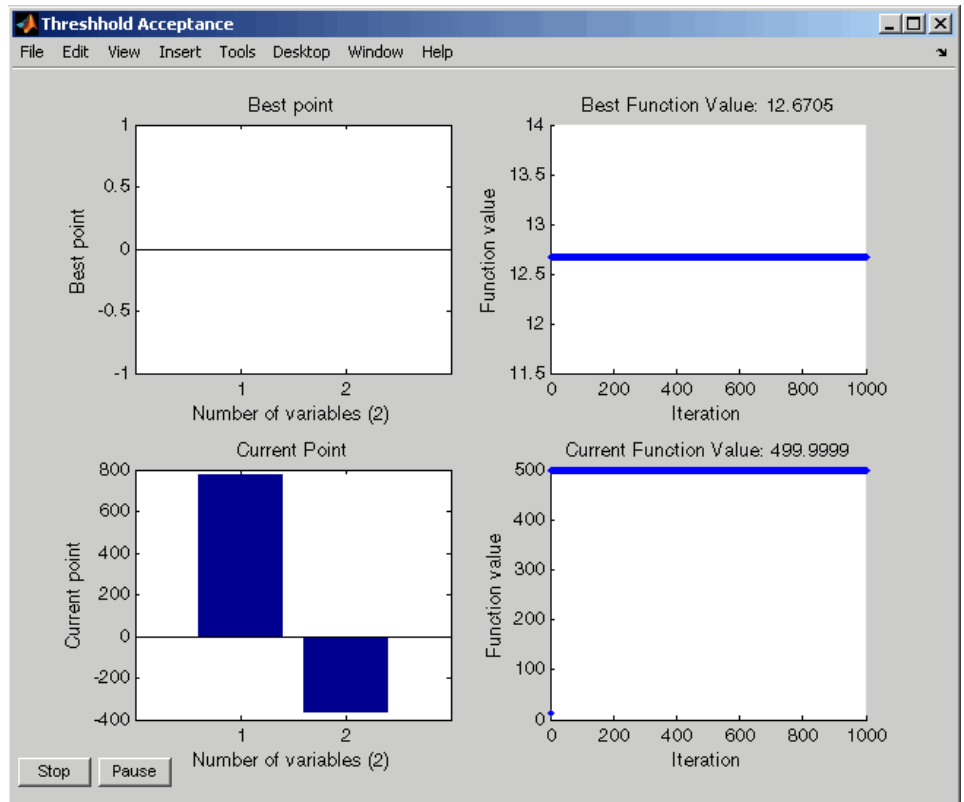
```
Optimization terminated: change in best function value  
less than options.TolFun.
```

```
x =  
0 0
```

```
fval =  
12.6705
```

```
exitflag =  
1
```

The plots displayed are shown below.



## See Also

saoptimset, threshacceptbnd, patternsearch, ga

# saoptimget

---

**Purpose** Values of simulated annealing or threshold acceptance algorithm options structure

**Syntax**  
`val = saoptimget(options, 'name')`  
`val = saoptimget(options, 'name', default)`

**Description** `val = saoptimget(options, 'name')` returns the value of the parameter name from the simulated annealing or threshold acceptance algorithm options structure options. `saoptimget(options, 'name')` returns an empty matrix [] if the value of name is not specified in options. It is only necessary to type enough leading characters of name to uniquely identify the parameter. `saoptimget` ignores case in parameter names.

`val = saoptimget(options, 'name', default)` returns the 'name' parameter, but returns the default value if the 'name' parameter is not specified (or is []) in options.

**Example**  
`opts = saoptimset('TolFun',1e-4);`  
`val = saoptimget(opts,'TolFun');`

returns `val = 1e-4` for TolFun.

**See Also** `simulannealbnd`, `threshacceptbnd`, `saoptimset`



|                    |   |
|--------------------|---|
| <b>Purpose</b>     | Create simulated annealing algorithm or threshold acceptance options structure  |
| <b>Syntax</b>      | <pre>saoptimset options = saoptimset options = saoptimset('param1',value1,'param2',value2,...) options = saoptimset(olddopts,'param1',value1,...) options = saoptimset(olddopts,newopts) options = saoptimset(optimfunction)</pre>  |
| <b>Description</b> | <p>saoptimset with no input or output arguments displays a complete list of parameters with their valid values.</p> <p>options = saoptimset (with no input arguments) creates a structure called options that contains the options, or <i>parameters</i>, for the simulated annealing or threshold acceptance algorithm with all parameters set to [].</p> <p>options = saoptimset('param1',value1,'param2',value2,...) creates a structure options and sets the value of 'param1' to value1, 'param2' to value2, and so on. Any unspecified parameters are set to []. It is sufficient to type only enough leading characters to define the parameter name uniquely. Case is ignored for parameter names. Note that for string values, correct case and the complete string are required.</p> <p>options = saoptimset(olddopts,'param1',value1,...) creates a copy of olddopts, modifying the specified parameters with the specified values.</p> <p>options = saoptimset(olddopts,newopts) combines an existing options structure, olddopts, with a new options structure, newopts. Any parameters in newopts with nonempty values overwrite the corresponding old parameters in olddopts.</p> <p>options = saoptimset(optimfunction) creates an options structure with all the parameter names and default values relevant to the optimization function optimfunction. optimfunction can be either 'simulannealbnd' or 'threshacceptbnd'. For example,</p> |

```
saoptimset('threshacceptbnd')

ans =
    AnnealingFcn: @annealingfast
    TemperatureFcn: @temperatureexp
    AcceptanceFcn: @acceptancethresh
        TolFun: 1.0000e-006
    StallIterLimit: '500*numberofvariables'
        MaxFunEvals: '3000*numberofvariables'
        TimeLimit: Inf
        MaxIter: Inf
    ObjectiveLimit: -Inf
        Display: 'final'
    DisplayInterval: 10
        HybridFcn: []
    HybridInterval: 'end'
        PlotFcns: []
    PlotInterval: 1
        OutputFcns: []
    InitialTemperature: 100
    ReannealInterval: 100
        DataType: 'double'
```

## Options

The following table lists the options you can set with `saoptimset`. See Chapter 9, “Options Reference” for a complete description of these options and their values. Values in {} denote the default value. You can also view the options parameters by typing `saoptimset` at the command line.

| Option             | Description  | Values  |
|--------------------|--|---|
| AcceptanceFcn      | Handle to the function the algorithm uses to determine if a new point is accepted                            | Function handle   @acceptancethresh   {@acceptancesa} for simulannealbnd or {@acceptancethresh} for threshacceptbnd |
| AnnealingFcn       | Handle to the function the algorithm uses to generate new points   | Function handle   @annealingboltz   {@annealingfast}  |
| DataType           | Type of decision variable  | 'custom'   {'double'}   |
| Display            | Level of display   | 'off'   'iter'   'diagnose'   {'final'}   |
| DisplayInterval    | Interval for iterative display   | Positive integer   {10}   |
| HybridFcn          | Automatically run HybridFcn (another optimization function) during or at the end of iterations of the solver | @fminsearch   @patternsearch   @fminunc   @fmincon   {}   |
| HybridInterval     | Interval (if not 'end' or 'never') at which HybridFcn is called  | Positive integer   'never'   {'end'}  |
| InitialTemperature | Initial value of temperature   | Positive integer   {100}  |
| MaxFunEvals        | Maximum number of objective function evaluations allowed   | Positive scalar   {3000*numberOfVariables}  |
| MaxIter            | Maximum number of iterations allowed   | Positive scalar   {Inf}   |

# saoptimset

| Option           | Description   | Values   |
|------------------|---|--|
| ObjectiveLimit   | Minimum objective function value desired  | Scalar   {Inf}   |
| OutputFcns       | Function(s) get(s) iterative data and can change options at run time  | Function handle or cell array of function handles   {[]}   |
| PlotFcns         | Plot function(s) called during iterations   | Function handle or cell array of function handles   @saplotbestf   @saplotbestx   @saplotf   @saplotstopping   @saplottemperature   {[]} |
| PlotInterval     | Plot functions are called at every interval   | Positive integer   {1}   |
| ReannealInterval | Reannealing interval  | Positive integer   {100}   |
| StallIterLimit   | Number of iterations over which average change in fitness function value at current point is less than options.TolFun | Positive integer   {500*numberOfVariables}   |
| TemperatureFcn   | Function used to update temperature schedule  | Function handle   @temperatureboltz   @temperaturefast   {@temperatureexp}   |
| TimeLimit        | The algorithm stops after running for TimeLimit seconds   | Positive scalar   {Inf}  |
| TolFun           | Termination tolerance on function value   | Positive scalar   {1e-6}   |

**See Also**      `simulannealbnd`, `threshacceptbnd`, `saoptimget`



## A

- accelerator
  - mesh 6-40
- algorithm
  - genetic 4-18
  - pattern search 3-14
  - simulated annealing 5-8
  - threshold acceptance 5-8
- annealing 5-6
- annealing schedule 5-6

## C

- cache 6-41
- children
  - crossover 4-20
  - elite 4-20
  - in genetic algorithms 4-17
  - mutation 4-20
- crossover 7-43
  - children 4-20
  - fraction 7-46

## D

- direct search 3-2
- diversity 4-16

## E

- elite children 4-20
- expansion
  - mesh 6-35
- exporting problems
  - from Genetic Algorithm Tool 7-16
  - from Pattern Search Tool 6-16

## F

- fitness function 4-15
  - vectorizing 7-60

- writing M-files 2-1
- fitness scaling 7-39

## G

- ga function 11-2
- gaoptimget function 11-8
- gaoptimset function 11-9
- gatool function 11-15
- generations 4-16
- genetic algorithm
  - description 4-18
  - nonlinear constraint algorithm, ALGA 4-25
  - options 9-23
  - overview 4-2
  - setting options at command line 7-26
  - stopping criteria 4-22
  - using from command line 7-24
- Genetic Algorithm Tool 7-2
  - defining problems 7-4
  - displaying plots 7-8
  - exporting options and problems 7-16
  - importing problems 7-19
  - opening 7-2
  - pausing and stopping 7-6
  - running 7-4
  - setting options 7-14
- global and local minima 7-52
- GPS 3-2

## H

- hybrid function 7-56

## I

- importing problems
  - to Genetic Algorithm Tool 7-19
  - to Pattern Search Tool 6-19
- individuals
  - applying the fitness function 4-15

initial population 4-19

## M

M-files

writing 2-2

MADS 3-2

maximizing functions 2-3

mesh 3-11

accelerator 6-40

expansion and contraction 6-35

minima

global 7-52

local 7-52

minimizing functions 2-3

mutation 7-43

options 7-44

## N

nonlinear constraint

genetic algorithm 7-7

pattern search 6-21

nonlinear constraint algorithms

ALGA 4-25

ALPS 3-23

## O

objective function 5-6

objective functions

writing M-files 2-1

options

genetic algorithm 9-23

simulated annealing algorithm 9-45

threshold acceptance algorithm 9-45

## P

parents in genetic algorithms 4-17

pattern search

description 3-14

nonlinear constraint algorithm, ALPS 3-23

options 9-2

overview 3-2

setting options at command line 6-22

terminology 3-10

using from command line 6-20

Pattern Search Tool 6-2

defining problems 6-4

displaying plots 6-9

exporting options and problems 6-16

importing problems 6-19

opening 6-3

pausing and stopping 6-9

running 6-5

setting options 6-15

patternsearch function 11-18

plots

genetic algorithm 7-8

pattern search 6-9

poll 3-12

complete 6-28

method 6-26

population 4-16

initial 4-19

initial range 7-35

options 7-34

size 7-38

psearchtool function 11-25

psoptimget function 11-27

psoptimset function 11-28

## R

Rastrigin's function 4-7

reannealing 5-6

reproduction options 7-43



**S**

- saoptimget function 11-42
- saoptimset function 11-43
- scaling
  - fitness 7-39
- search method 6-32
- selection function 7-42
- setting options
  - genetic algorithm 7-34
  - pattern search 6-26
  - simulated annealing algorithm 8-7
  - threshold acceptance algorithm 8-7
- simulannealbnd function 11-32
- simulated annealing
  - description 5-8
  - overview 5-2
- simulated annealing algorithm
  - options 9-45
  - setting options at command line 8-3
  - stopping criteria 5-8

- using from command line 8-2

- stopping criteria
  - pattern search 3-21

**T**

- temperature 5-6
- threshacceptbnd function 11-37
- threshold acceptance
  - description 5-8
  - overview 5-2
- threshold acceptance algorithm
  - options 9-45
  - setting options at command line 8-3
  - stopping criteria 5-8
  - using from command line 8-2

**V**

- vectorizing fitness functions 7-60