


Optimization Toolbox™

User's Guide

R2014a



MATLAB®

How to Contact MathWorks



www.mathworks.com Web
comp.soft-sys.matlab Newsgroup
www.mathworks.com/contact_TS.html Technical Support



suggest@mathworks.com Product enhancement suggestions
bugs@mathworks.com Bug reports
doc@mathworks.com Documentation error reports
service@mathworks.com Order status, license renewals, passcodes
info@mathworks.com 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.

Optimization Toolbox™ User's Guide

© COPYRIGHT 1990–2014 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 and Simulink are registered trademarks of The MathWorks, Inc. See www.mathworks.com/trademarks for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

Patents

MathWorks products are protected by one or more U.S. patents. Please see www.mathworks.com/patents for more information.

Revision History

November 1990	First printing	
December 1996	Second printing	For MATLAB® 5
January 1999	Third printing	For Version 2 (Release 11)
September 2000	Fourth printing	For Version 2.1 (Release 12)
June 2001	Online only	Revised for Version 2.1.1 (Release 12.1)
September 2003	Online only	Revised for Version 2.3 (Release 13SP1)
June 2004	Fifth printing	Revised for Version 3.0 (Release 14)
October 2004	Online only	Revised for Version 3.0.1 (Release 14SP1)
March 2005	Online only	Revised for Version 3.0.2 (Release 14SP2)
September 2005	Online only	Revised for Version 3.0.3 (Release 14SP3)
March 2006	Online only	Revised for Version 3.0.4 (Release 2006a)
September 2006	Sixth printing	Revised for Version 3.1 (Release 2006b)
March 2007	Seventh printing	Revised for Version 3.1.1 (Release 2007a)
September 2007	Eighth printing	Revised for Version 3.1.2 (Release 2007b)
March 2008	Online only	Revised for Version 4.0 (Release 2008a)
October 2008	Online only	Revised for Version 4.1 (Release 2008b)
March 2009	Online only	Revised for Version 4.2 (Release 2009a)
September 2009	Online only	Revised for Version 4.3 (Release 2009b)
March 2010	Online only	Revised for Version 5.0 (Release 2010a)
September 2010	Online only	Revised for Version 5.1 (Release 2010b)
April 2011	Online only	Revised for Version 6.0 (Release 2011a)
September 2011	Online only	Revised for Version 6.1 (Release 2011b)
March 2012	Online only	Revised for Version 6.2 (Release 2012a)
September 2012	Online only	Revised for Version 6.2.1 (Release 2012b)
March 2013	Online only	Revised for Version 6.3 (Release 2013a)
September 2013	Online only	Revised for Version 6.4 (Release 2013b)
March 2014	Online only	Revised for Version 7.0 (Release 2014a)

Acknowledgments

MathWorks would like to acknowledge the following contributors to Optimization Toolbox™ algorithms.

Thomas F. Coleman researched and contributed algorithms for constrained and unconstrained minimization, nonlinear least squares and curve fitting, constrained linear least squares, quadratic programming, and nonlinear equations.

Dr. Coleman is Professor of Combinatorics and Optimization at the University of Waterloo.

Yin Zhang researched and contributed the large-scale linear programming algorithm.

Dr. Zhang is Professor of Computational and Applied Mathematics at Rice University.

Acknowledgments

Getting Started

1

Optimization Toolbox Product Description	1-2
Key Features	1-2
Solve a Constrained Nonlinear Problem	1-3
Problem Formulation: Rosenbrock's Function	1-3
Defining the Problem in Toolbox Syntax	1-4
Running the Optimization	1-6
Interpreting the Result	1-11
Set Up a Linear Program	1-13
Convert a Problem to Solver Form	1-13
Model Description	1-13
Solution Method	1-15
Bibliography	1-22

Setting Up an Optimization

2

Optimization Toolbox Solvers	2-2
Choosing a Solver	2-4
Optimization Decision Table	2-4
Choosing the Algorithm	2-7
Problems Handled by Optimization Toolbox Functions ...	2-14
Writing Objective Functions	2-18
Types of Objective Functions	2-18
Writing Scalar Objective Functions	2-19
Writing Vector and Matrix Objective Functions	2-27

Writing Objective Functions for Linear or Quadratic Problems	2-30
Maximizing an Objective	2-31
Matrix Arguments	2-32
Writing Constraints	2-33
Types of Constraints	2-33
Iterations Can Violate Constraints	2-34
Bound Constraints	2-35
Linear Inequality Constraints	2-36
Linear Equality Constraints	2-37
Nonlinear Constraints	2-37
Or Instead of And Constraints	2-41
How to Use All Types of Constraints	2-46
Objective and Nonlinear Constraints in the Same Function	2-48
Passing Extra Parameters	2-53
Extra Parameters, Fixed Variables, or Data	2-53
Anonymous Functions	2-53
Nested Functions	2-55
Global Variables	2-56
What Are Options?	2-57
Options in Common Use: Tuning and Troubleshooting	2-58
Set and Change Options	2-59
Choose Between <code>optimoptions</code> and <code>optimset</code>	2-60
View Options	2-63
Tolerances and Stopping Criteria	2-65
Checking Validity of Gradients or Jacobians	2-69

How to Check Derivatives	2-69
Example: Checking Derivatives of Objective and Constraint Functions	2-70

Examining Results

3

Current Point and Function Value	3-2
Exit Flags and Exit Messages	3-3
Exit Flags	3-3
Exit Messages	3-5
Enhanced Exit Messages	3-5
Exit Message Options	3-9
Iterations and Function Counts	3-11
First-Order Optimality Measure	3-12
What Is First-Order Optimality Measure?	3-12
Stopping Rules Related to First-Order Optimality	3-12
Unconstrained Optimality	3-13
Constrained Optimality Theory	3-13
Constrained Optimality in Solver Form	3-15
Iterative Display	3-17
Introduction	3-17
Common Headings	3-18
Function-Specific Headings	3-18
Output Structures	3-26
Lagrange Multiplier Structures	3-27
Hessian	3-28
fminunc Hessian	3-28
fmincon Hessian	3-29

Plot Functions	3-31
Plot an Optimization During Execution	3-31
Using a Plot Function	3-31
Output Functions	3-37
What Is an Output Function?	3-37
Example: Using Output Functions	3-37

Steps to Take After Running a Solver

4

Overview of Next Steps	4-2
When the Solver Fails	4-3
Too Many Iterations or Function Evaluations	4-3
No Feasible Point	4-7
Problem Unbounded	4-9
fsolve Could Not Solve Equation	4-10
Solver Takes Too Long	4-10
When the Solver Might Have Succeeded	4-15
Final Point Equals Initial Point	4-15
Local Minimum Possible	4-15
When the Solver Succeeds	4-23
What Can Be Wrong If The Solver Succeeds?	4-23
1. Change the Initial Point	4-24
2. Check Nearby Points	4-25
3. Check your Objective and Constraint Functions	4-26
Local vs. Global Optima	4-27
Optimizing a Simulation or Ordinary Differential Equation	4-33
What Is Optimizing a Simulation or ODE?	4-33
Potential Problems and Solutions	4-33
Bibliography	4-39

5

Optimization App	5-2
Optimization App Basics	5-2
Specifying Certain Options	5-10
Importing and Exporting Your Work	5-13

Optimization Algorithms and Examples

6

Optimization Theory Overview	6-4
Unconstrained Nonlinear Optimization Algorithms ...	6-5
Unconstrained Optimization Definition	6-5
fminunc trust-region Algorithm	6-5
fminunc quasi-newton Algorithm	6-8
fminsearch Algorithm	6-14
fminunc Unconstrained Minimization	6-17
Step 1: Write a file objfun.m.	6-17
Step 2: Invoke one of the unconstrained optimization routines.	6-17
Minimization with Gradient and Hessian	6-20
Step 1: Write a file brownfgh.m that computes the objective function, the gradient of the objective, and the sparse tridiagonal Hessian matrix.	6-20
Step 2: Call a nonlinear minimization routine with a starting point xstart.	6-20
Minimization with Gradient and Hessian Sparsity Pattern	6-22
Step 1: Write a file brownfg.m that computes the objective function and the gradient of the objective.	6-22
Step 2: Call a nonlinear minimization routine with a starting point xstart.	6-23

Constrained Nonlinear Optimization Algorithms	6-26
Constrained Optimization Definition	6-26
fmincon Trust Region Reflective Algorithm	6-26
fmincon Active Set Algorithm	6-32
fmincon SQP Algorithm	6-42
fmincon Interior Point Algorithm	6-43
fminbnd Algorithm	6-47
fseminf Problem Formulation and Algorithm	6-47
Nonlinear Inequality Constraints	6-52
Step 1: Write a file objfun.m for the objective function. ...	6-52
Step 2: Write a file confun.m for the constraints.	6-52
Step 3: Invoke constrained optimization routine.	6-53
Nonlinear Constraints with Gradients	6-54
Step 1: Write a file for the objective function and gradient.	6-54
Step 2: Write a file for the nonlinear constraints and the gradients of the nonlinear constraints.	6-54
Step 3: Invoke the constrained optimization routine.	6-55
fmincon Interior-Point Algorithm with Analytic Hessian	6-57
Nonlinear Equality and Inequality Constraints	6-65
Step 1: Write a file objfun.m.	6-65
Step 2: Write a file confuneq.m for the nonlinear constraints.	6-65
Step 3: Invoke constrained optimization routine.	6-65
Optimization App with the fmincon Solver	6-67
Step 1: Write a file objecfun.m for the objective function. ..	6-67
Step 2: Write a file nonlconstr.m for the nonlinear constraints.	6-67
Step 3: Set up and run the problem with the Optimization app.	6-68
Minimization with Bound Constraints and Banded Preconditioner	6-72
Step 1: Write a file tbroyfg.m that computes the objective function and the gradient of the objective	6-72

Step 2: Call a nonlinear minimization routine with a starting point xstart.	6-74
Minimization with Linear Equality Constraints	6-77
Step 1: Write a file brownfgh.m that computes the objective function, the gradient of the objective, and the sparse tridiagonal Hessian matrix.	6-77
Step 2: Call a nonlinear minimization routine with a starting point xstart.	6-78
Minimization with Dense Structured Hessian, Linear Equalities	6-80
Hessian Multiply Function for Lower Memory	6-80
Step 1: Write a file brownvv.m that computes the objective function, the gradient, and the sparse part of the Hessian.	6-81
Step 2: Write a function to compute Hessian-matrix products for H given a matrix Y.	6-81
Step 3: Call a nonlinear minimization routine with a starting point and linear equality constraints.	6-82
Preconditioning	6-84
Symbolic Math Toolbox Calculates Gradients and Hessians	6-85
Create the Variables	6-87
Include the Linear Constraints	6-88
Create the Nonlinear Constraints, Their Gradients and Hessians	6-90
Create the Objective Function, Its Gradient and Hessian	6-91
Create the Objective Function File	6-92
Create the Constraint Function File	6-92
Generate the Hessian Files	6-93
Run the Optimization	6-94
Compare to Optimization Without Gradients and Hessians	6-97
Clear the Symbolic Variable Assumptions	6-99
One-Dimensional Semi-Infinite Constraints	6-100
Two-Dimensional Semi-Infinite Constraint	6-103

Linear Programming Algorithms	6-106
Linear Programming Definition	6-106
Interior-Point Linear Programming	6-106
Active-Set linprog Algorithm	6-110
linprog Simplex Algorithm	6-114
Typical Linear Programming Problem	6-119
Quadratic Programming Algorithms	6-121
Quadratic Programming Definition	6-121
interior-point-convex quadprog Algorithm	6-121
trust-region-reflective quadprog Algorithm	6-126
active-set quadprog Algorithm	6-131
Quadratic Minimization with Bound Constraints	6-135
Step 1: Load the Hessian and define f, lb, and ub.	6-135
Step 2: Call a quadratic minimization routine with a starting point xstart.	6-135
Quadratic Minimization with Dense, Structured Hessian	6-138
Step 1: Decide what part of H to pass to quadprog as the first argument.	6-138
Step 2: Write a function to compute Hessian-matrix products for H.	6-138
Step 3: Call a quadratic minimization routine with a starting point.	6-140
Preconditioning	6-141
Large Sparse Quadratic Program with Interior Point Algorithm	6-144
Binary Integer Programming Algorithms	6-148
Binary Integer Programming Definition	6-148
bintprog Algorithm	6-148
Optimal Investments Via Binary Integer Programming	6-151
Investments with Constraints	6-151
Problem Statement	6-151
bintprog Formulation	6-153

bintprog Solution	6-155
Mixed-Integer Linear Programming Algorithms	6-157
Mixed-Integer Linear Programming Definition	6-157
intlinprog Algorithm	6-157
Tuning Integer Linear Programming	6-166
Change Options to Improve the Solution Process	6-166
Some “Integer” Solutions Are Not Integers	6-167
Large Components Not Integer Valued	6-168
Large Coefficients Disallowed	6-168
Mixed-Integer Linear Programming Basics	6-169
Factory, Warehouse, Sales Allocation Model	6-173
Travelling Salesman Problem	6-185
Solve Sudoku Puzzles Via Integer Programming	6-192
Least-Squares (Model Fitting) Algorithms	6-200
Least Squares Definition	6-200
Trust-Region-Reflective Least Squares	6-201
Levenberg-Marquardt Method	6-205
lsqnonlin with a Simulink Model	6-209
Nonlinear Least Squares With and Without Jacobian ..	6-215
Jacobians and Memory in Nonlinear Least Squares	6-215
Step 1: Write a file myfun.m that computes the objective function values.	6-216
Step 2: Call the nonlinear least-squares routine.	6-216
Step 3: Include a Jacobian.	6-217
Linear Least Squares with Bound Constraints	6-219
Optimization App with the lsqin Solver	6-221
The Problem	6-221
Setting Up the Problem	6-221

Jacobian Multiply Function with Linear Least Squares	6-224
Nonlinear Curve Fitting with lsqcurvefit	6-229
Fit a Model to Complex-Valued Data	6-231
Multiobjective Optimization Algorithms	6-236
Multiobjective Optimization Definition	6-236
Algorithms	6-237
Using fminimax with a Simulink Model	6-242
Signal Processing Using fgoalattain	6-246
Step 1: Write a file filtmin.m	6-247
Step 2: Invoke optimization routine	6-247
Equation Solving Algorithms	6-249
Equation Solving Definition	6-249
Trust-Region Reflective fsolve Algorithm	6-249
Trust-Region Dogleg Method	6-252
Levenberg-Marquardt Method	6-255
\ Algorithm	6-255
fzero Algorithm	6-255
Nonlinear Equations with Analytic Jacobian	6-257
Step 1: Write a file bananaobj.m to compute the objective function values and the Jacobian.	6-258
Step 2: Call the solve routine for the system of equations.	6-259
Nonlinear Equations with Finite-Difference Jacobian	6-261
Nonlinear Equations with Jacobian	6-263
Step 1: Write a file nlsf1.m that computes the objective function values and the Jacobian.	6-263
Step 2: Call the solve routine for the system of equations.	6-264

Nonlinear Equations with Jacobian Sparsity	
Pattern	6-266
Step 1: Write a file nlsf1a.m that computes the objective function values.	6-266
Step 2: Call the system of equations solve routine.	6-267
Nonlinear Systems with Constraints	6-269
Solve Equations with Inequality Constraints	6-269
Use Different Start Points	6-270
Use Different Algorithms	6-270
Use lsqnonlin with Bounds	6-271
Set Equations and Inequalities as fmincon Constraints ..	6-272
Bibliography	6-274

Parallel Computing for Optimization

7

Parallel Computing in Optimization Toolbox	
Functions	7-2
Parallel Optimization Functionality	7-2
Parallel Estimation of Gradients	7-3
Nested Parallel Functions	7-4
Using Parallel Computing in Optimization Toolbox ...	7-5
Using Parallel Computing with Multicore Processors	7-5
Using Parallel Computing with a Multiprocessor Network	7-6
Testing Parallel Computations	7-7
Minimizing an Expensive Optimization Problem Using Parallel Computing Toolbox™	7-9
Improving Performance with Parallel Computing	7-16
Factors That Affect Speed	7-16
Factors That Affect Results	7-17
Searching for Global Optima	7-17

External Interface

8

ktrlink: An Interface to KNITRO Libraries	8-2
What Is ktrlink?	8-2
Installation and Configuration	8-2
Using ktrlink to Optimize	8-4
Setting Options	8-8
Sparse Matrix Considerations	8-9

Argument and Options Reference

9

Function Arguments	9-2
Input Arguments	9-2
Output Arguments	9-5
Optimization Options Reference	9-7
Optimization Options	9-7
Output Function	9-21
Plot Functions	9-30

Functions — Alphabetical List

10

Getting Started

- “Optimization Toolbox Product Description” on page 1-2
- “Solve a Constrained Nonlinear Problem” on page 1-3
- “Set Up a Linear Program” on page 1-13

Optimization Toolbox Product Description

Solve linear, quadratic, integer, and nonlinear optimization problems

Optimization Toolbox provides functions for finding parameters that minimize or maximize objectives while satisfying constraints. The toolbox includes solvers for linear programming, mixed-integer linear programming, quadratic programming, nonlinear optimization, and nonlinear least squares. You can use these solvers to find optimal solutions to continuous and discrete problems, perform tradeoff analyses, and incorporate optimization methods into algorithms and applications.

Key Features

- Nonlinear and multiobjective optimization
- Solvers for nonlinear least squares, data fitting, and nonlinear equations
- Quadratic and linear programming
- Mixed-integer linear programming
- Optimization app for defining and solving optimization problems and monitoring solution progress
- Acceleration of constrained nonlinear solvers with Parallel Computing Toolbox™

Solve a Constrained Nonlinear Problem

In this section...

“Problem Formulation: Rosenbrock’s Function” on page 1-3

“Defining the Problem in Toolbox Syntax” on page 1-4

“Running the Optimization” on page 1-6

“Interpreting the Result” on page 1-11

Problem Formulation: Rosenbrock’s Function

Consider the problem of minimizing Rosenbrock’s function


$$f(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2,$$

over the *unit disk*, i.e., the disk of radius 1 centered at the origin. In other

words, find x that minimizes the function $f(x)$ over the set $x_1^2 + x_2^2 \leq 1$. This problem is a minimization of a nonlinear function with a nonlinear constraint.

Note Rosenbrock’s function is a standard test function in optimization. It has a unique minimum value of 0 attained at the point (1,1). Finding the minimum is a challenge for some algorithms since it has a shallow minimum inside a deeply curved valley.

Here are two views of Rosenbrock’s function in the unit disk. The vertical axis is log-scaled; in other words, the plot shows $\log(1+f(x))$. Contour lines lie beneath the surface plot.



Rosenbrock's function, log-scaled: two views.

The function $f(x)$ is called the *objective function*. This is the function you wish to minimize. The inequality $x_1^2 + x_2^2 \leq 1$ is called a *constraint*. Constraints limit the set of x over which you may search for a minimum. You can have any number of constraints, which are inequalities or equations.

All Optimization Toolbox optimization functions minimize an objective function. To maximize a function f , apply an optimization routine to minimize $-f$. For more details about maximizing, see “Maximizing an Objective” on page 2-31.

Defining the Problem in Toolbox Syntax

To use Optimization Toolbox software, you need to

- 1 Define your objective function in the MATLAB® language, as a function file or anonymous function. This example will use a function file.
- 2 Define your constraint(s) as a separate file or anonymous function.

Function File for Objective Function

A function file is a text file containing MATLAB commands with the extension .m. Create a new function file in any text editor, or use the built-in MATLAB Editor as follows:

- 1 At the command line enter:

```
edit rosenbrock
```

The MATLAB Editor opens.

- 2 In the editor enter:

```
function f = rosenbrock(x)  
f = 100*(x(2) - x(1)^2)^2 + (1 - x(1))^2;
```

- 3 Save the file by selecting **File > Save**.

File for Constraint Function

Constraint functions must be formulated so that they are in the form

$c(x) \leq 0$ or $ceq(x) = 0$. The constraint $x_1^2 + x_2^2 \leq 1$ needs to be reformulated as $x_1^2 + x_2^2 - 1 \leq 0$ in order to have the correct syntax.

Furthermore, toolbox functions that accept nonlinear constraints need to have both equality and inequality constraints defined. In this example there is only an inequality constraint, so you must pass an empty array [] as the equality constraint function *ceq*.

With these considerations in mind, write a function file for the nonlinear constraint:

- 1 Create a file named `unitdisk.m` containing the following code:

```
function [c, ceq] = unitdisk(x)
c = x(1)^2 + x(2)^2 - 1;
ceq = [ ];
```

2 Save the file `unitdisk.m`.


Running the Optimization

There are two ways to run the optimization:

- Using the “Optimization app” on page 1-6
- Using command line functions; see “Minimizing at the Command Line” on page 1-10.

Optimization app

1 Start the Optimization app by typing `optimtool` at the command line.



For more information about this tool, see “Optimization App” on page 5-2.

- The default **Solver** fmincon - Constrained nonlinear minimization is selected. This solver is appropriate for this problem, since Rosenbrock’s function is nonlinear, and the problem has a constraint. For more information about how to choose a solver, see “Choosing a Solver” on page 2-4.
- In the **Algorithm** pop-up menu choose Interior point, which is the default.

- 4 For **Objective function** enter @rosenbrock. The @ character indicates that this is a function handle of the file rosenbrock.m.
- 5 For **Start point** enter [0 0]. This is the initial point where fmincon begins its search for a minimum.
- 6 For **Nonlinear constraint function** enter @unitdisk, the function handle of unitdisk.m.

Your **Problem Setup and Results** pane should match this figure.

Problem Setup and Results

Solver:

Algorithm:

Problem

Objective function:

Derivatives:

Start point:

Constraints:

Linear inequalities: A: b:


Linear equalities: Aeq: beq:

Bounds: Lower: Upper:


Nonlinear constraint function:

Derivatives:

- 7 In the Options pane (center bottom), select iterative in the **Level of display** pop-up menu. (If you don't see the option, click **Display to command window**.) This shows the progress of fmincon in the command window.



8 Click **Start** under **Run solver and view results**.



The following message appears in the box below the **Start** button:

```
Optimization running.
Objective function value: 0.045674824758137236
Local minimum found that satisfies the constraints.
```

```
Optimization completed because the objective function is non-decreasing in
feasible directions, to within the default value of the function tolerance,
and constraints are satisfied to within the default value of the constraint tolerance.
```

Your objective function value may differ slightly, depending on your computer system and version of Optimization Toolbox software.

The message tells you that:

- The search for a constrained optimum ended because the derivative of the objective function is nearly 0 in directions allowed by the constraint.
- The constraint is satisfied to the requisite accuracy.

“Exit Flags and Exit Messages” on page 3-3 discusses exit messages such as these.

The minimizer x appears under **Final point**.

Final point:	
1 ▲	2
0.786	0.618

Minimizing at the Command Line

You can run the same optimization from the command line, as follows.

- 1 Create an options structure to choose iterative display and the interior-point algorithm:

```
options = optimoptions(@fmincon,...
    'Display','iter','Algorithm','interior-point');
```

- 2 Run the `fmincon` solver with the `options` structure, reporting both the location `x` of the minimizer, and value `fval` attained by the objective function:

```
[x,fval] = fmincon(@rosenbrock,[0 0],...
    [],[],[],[],[],[],@unitdisk,options)
```

The six sets of empty brackets represent optional constraints that are not being used in this example. See the `fmincon` function reference pages for the syntax.

MATLAB outputs a table of iterations, and the results of the optimization:

```
Local minimum found that satisfies the constraints.
```

```
Optimization completed because the objective function is non-decreasing in
feasible directions, to within the selected value of the function tolerance,
and constraints are satisfied to within the selected value of the constraint tolerance.
```

```
x =
    0.7864    0.6177
```

```
fval =
    0.0457
```

The message tells you that the search for a constrained optimum ended because the derivative of the objective function is nearly 0 in directions allowed by the constraint, and that the constraint is satisfied to the requisite accuracy. Several phrases in the message contain links that give you more information about the terms used in the message. For more details about these links, see “Enhanced Exit Messages” on page 3-5.

Interpreting the Result

The iteration table in the command window shows how MATLAB searched for the minimum value of Rosenbrock’s function in the unit disk. This table is the same whether you use Optimization app or the command line. MATLAB reports the minimization as follows:

Iter	F-count	f(x)	Feasibility	First-order optimality	Norm of step
0	3	1.000000e+00	0.000e+00	2.000e+00	
1	13	7.753537e-01	0.000e+00	6.250e+00	1.768e-01
2	18	6.519648e-01	0.000e+00	9.048e+00	1.679e-01
3	21	5.543209e-01	0.000e+00	8.033e+00	1.203e-01
4	24	2.985207e-01	0.000e+00	1.790e+00	9.328e-02
5	27	2.653799e-01	0.000e+00	2.788e+00	5.723e-02
6	30	1.897216e-01	0.000e+00	2.311e+00	1.147e-01
7	33	1.513701e-01	0.000e+00	9.706e-01	5.764e-02
8	36	1.153330e-01	0.000e+00	1.127e+00	8.169e-02
9	39	1.198058e-01	0.000e+00	1.000e-01	1.522e-02
10	42	8.910052e-02	0.000e+00	8.378e-01	8.301e-02
11	45	6.771960e-02	0.000e+00	1.365e+00	7.149e-02
12	48	6.437664e-02	0.000e+00	1.146e-01	5.701e-03
13	51	6.329037e-02	0.000e+00	1.883e-02	3.774e-03
14	54	5.161934e-02	0.000e+00	3.016e-01	4.464e-02
15	57	4.964194e-02	0.000e+00	7.913e-02	7.894e-03
16	60	4.955404e-02	0.000e+00	5.462e-03	4.185e-04
17	63	4.954839e-02	0.000e+00	3.993e-03	2.208e-05
18	66	4.658289e-02	0.000e+00	1.318e-02	1.255e-02
19	69	4.647011e-02	0.000e+00	8.006e-04	4.940e-04
20	72	4.569141e-02	0.000e+00	3.136e-03	3.379e-03

21	75	4.568281e-02	0.000e+00	6.439e-05	3.974e-05
22	78	4.568281e-02	0.000e+00	8.000e-06	1.083e-07
23	81	4.567641e-02	0.000e+00	1.601e-06	2.793e-05
24	84	4.567482e-02	0.000e+00	2.062e-08	6.916e-06

This table might differ from yours depending on toolbox version and computing platform. The following description applies to the table as displayed.

- The first column, labeled `Iter`, is the iteration number from 0 to 24. `fmincon` took 24 iterations to converge.
- The second column, labeled `F-count`, reports the cumulative number of times Rosenbrock's function was evaluated. The final row shows an `F-count` of 84, indicating that `fmincon` evaluated Rosenbrock's function 84 times in the process of finding a minimum.
- The third column, labeled `f(x)`, displays the value of the objective function. The final value, 0.04567482, is the minimum that is reported in the Optimization app **Run solver and view results** box, and at the end of the exit message in the command window.
- The fourth column, `Feasibility`, is 0 for all iterations. This column shows the value of the constraint function `unitdisk` at each iteration where the constraint is positive. Since the value of `unitdisk` was negative in all iterations, every iteration satisfied the constraint.

The other columns of the iteration table are described in "Iterative Display" on page 3-17.

Set Up a Linear Program

In this section...

“Convert a Problem to Solver Form” on page 1-13

“Model Description” on page 1-13

“Solution Method” on page 1-15

“Bibliography” on page 1-22

Convert a Problem to Solver Form

This example shows how to convert a problem from mathematical form into Optimization Toolbox solver syntax. While the problem is a linear program, the techniques apply to all solvers.

The variables and expressions in the problem represent a model of operating a chemical plant, from an example in Edgar and Himmelblau [1]. There are two videos that describe the problem.

- Optimization Modeling 1 shows the problem in pictorial form. It shows how to generate the mathematical expressions of “Model Description” on page 1-13 from the picture.
- Optimization Modeling 2 describes how to convert these mathematical expressions into Optimization Toolbox solver syntax. This video shows how to solve the problem, and how to interpret the results.

The remainder of this example is concerned solely with transforming the problem to solver syntax. The example closely follows the video Optimization Modeling 2. The main difference between the video and the example is that this example shows how to use named variables, or index variables, which are similar to hash keys. This difference is in “Combine Variables Into One Vector” on page 1-16.

Model Description

The video Optimization Modeling 1 suggests that one way to convert a problem into mathematical form is to:

- 1** Get an overall idea of the problem
- 2** Identify the goal (maximizing or minimizing something)
- 3** Identify (name) variables
- 4** Identify constraints
- 5** Determine which variables you can control
- 6** Specify all quantities in mathematical notation
- 7** Check the model for completeness and correctness

For the meaning of the variables in this section, see the video Optimization Modeling 1.

The optimization problem is to minimize the objective function, subject to all the other expressions as constraints.

The objective function is:

$$0.002614 \text{ HPS} + 0.0239 \text{ PP} + 0.009825 \text{ EP}.$$

The constraints are:

$$2500 \leq P1 \leq 6250$$

$$I1 \leq 192,000$$

$$C \leq 62,000$$

$$I1 - HE1 \leq 132,000$$

$$I1 = LE1 + HE1 + C$$

$$1359.8 I1 = 1267.8 HE1 + 1251.4 LE1 + 192 C + 3413 P1$$

$$3000 \leq P2 \leq 9000$$

$$I2 \leq 244,000$$

$$LE2 \leq 142,000$$

$$I2 = LE2 + HE2$$

$$1359.8 I2 = 1267.8 HE2 + 1251.4 LE2 + 3413 P2$$

$$\text{HPS} = I1 + I2 + \text{BF1}$$

$$\text{HPS} = C + \text{MPS} + \text{LPS}$$

$$\text{LPS} = LE1 + LE2 + \text{BF2}$$

$$\text{MPS} = HE1 + HE2 + \text{BF1} - \text{BF2}$$

$$P1 + P2 + PP \geq 24,550$$

$$EP + PP \geq 12,000$$

$$MPS \geq 271,536$$

$$LPS \geq 100,623$$

All variables are positive.

Solution Method

To solve the optimization problem, take the following steps.

- 1 “Choose a Solver” on page 1-15
- 2 “Combine Variables Into One Vector” on page 1-16
- 3 “Write Bound Constraints” on page 1-18
- 4 “Write Linear Inequality Constraints” on page 1-18
- 5 “Write Linear Equality Constraints” on page 1-19
- 6 “Write the Objective” on page 1-20
- 7 “Solve the Problem with linprog” on page 1-20
- 8 “Examine the Solution” on page 1-21

The steps are also shown in the video Optimization Modeling 2.

Choose a Solver

To find the appropriate solver for this problem, consult the “Optimization Decision Table” on page 2-4. The table asks you to categorize your problem by type of objective function and types of constraints. For this problem, the objective function is linear, and the constraints are linear. The decision table recommends using the linprog solver.

As you see in “Problems Handled by Optimization Toolbox Functions” on page 2-14 or the linprog function reference page, the linprog solver solves problems of the form

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ lb \leq x \leq ub. \end{cases} \quad (1-1)$$

- $f^T x$ means a row vector of constants f multiplying a column vector of variables x . In other words,

$$f^T x = f(1)x(1) + f(2)x(2) + \dots + f(n)x(n),$$

where n is the length of f .

- $A x \leq b$ represents linear inequalities. A is a k -by- n matrix, where k is the number of inequalities and n is the number of variables (size of x). b is a vector of length k . For more information, see “Linear Inequality Constraints” on page 2-36.
- $Aeq x = beq$ represents linear equalities. Aeq is an m -by- n matrix, where m is the number of equalities and n is the number of variables (size of x). beq is a vector of length m . For more information, see “Linear Equality Constraints” on page 2-37.
- $lb \leq x \leq ub$ means each element in the vector x must be greater than the corresponding element of lb , and must be smaller than the corresponding element of ub . For more information, see “Bound Constraints” on page 2-35.

The syntax of the `linprog` solver, as shown in its function reference page, is

```
[x fval] = linprog(f,A,b,Aeq,beq,lb,ub);
```

The inputs to the `linprog` solver are the matrices and vectors in Equation 1-1.

Combine Variables Into One Vector

There are 16 variables in the equations of “Model Description” on page 1-13. Put these variables into one vector. The name of the vector of variables is x in Equation 1-1. Decide on an order, and construct the components of x out of the variables.

The following code constructs the vector using a cell array of strings. Each string is the name of a variable.

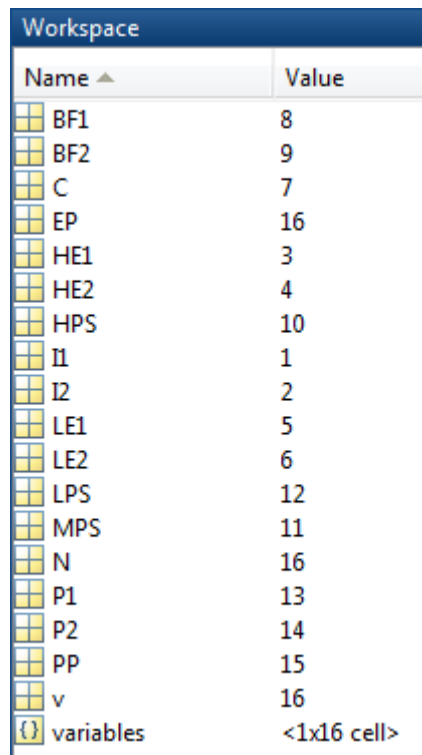
```
variables = {'I1','I2','HE1','HE2','LE1','LE2','C','BF1',...
```

```

    'BF2','HPS','MPS','LPS','P1','P2','PP','EP'];
N = length(variables);
% create variables for indexing
for v = 1:N
    eval([variables{v},' = ', num2str(v),';']);
end

```

Executing these commands creates the following named variables in your workspace:



Workspace	
Name ▲	Value
BF1	8
BF2	9
C	7
EP	16
HE1	3
HE2	4
HPS	10
I1	1
I2	2
LE1	5
LE2	6
LPS	12
MPS	11
N	16
P1	13
P2	14
PP	15
v	16
variables	<1x16 cell>

These named variables represent index numbers for the components of x . You do not have to create named variables. The video Optimization Modeling 2 shows how to solve the problem simply using the index numbers of the components of x .

Write Bound Constraints

There are four variables with lower bounds, and six with upper bounds in the equations of “Model Description” on page 1-13. The lower bounds:

$$\begin{aligned}P1 &\geq 2500 \\P2 &\geq 3000 \\MPS &\geq 271,536 \\LPS &\geq 100,623.\end{aligned}$$

Also, all the variables are positive, which means they have a lower bound of zero.

Create the lower bound vector `lb` as a vector of 0, then add the four other lower bounds.

```
lb = zeros(size(variables));  
lb([P1,P2,MPS,LPS]) = ...  
    [2500,3000,271536,100623];
```

The variables with upper bounds are:

$$\begin{aligned}P1 &\leq 6250 \\P2 &\leq 9000 \\I1 &\leq 192,000 \\I2 &\leq 244,000 \\C &\leq 62,000 \\LE2 &\leq 142000.\end{aligned}$$

Create the upper bound vector as a vector of `Inf`, then add the six upper bounds.

```
ub = Inf(size(variables));  
ub([P1,P2,I1,I2,C,LE2]) = ...  
    [6250,9000,192000,244000,62000,142000];
```

Write Linear Inequality Constraints

There are three linear inequalities in the equations of “Model Description” on page 1-13:

$$\begin{aligned} I1 - HE1 &\leq 132,000 \\ EP + PP &\geq 12,000 \\ P1 + P2 + PP &\geq 24,550. \end{aligned}$$

In order to have the equations in the form $Ax \leq b$, put all the variables on the left side of the inequality. All these equations already have that form. Ensure that each inequality is in “less than” form by multiplying through by -1 wherever appropriate:

$$\begin{aligned} I1 - HE1 &\leq 132,000 \\ -EP - PP &\leq -12,000 \\ -P1 - P2 - PP &\leq -24,550. \end{aligned}$$

In your MATLAB workspace, create the A matrix as a 3-by-16 zero matrix, corresponding to 3 linear inequalities in 16 variables. Create the b vector with three components.

```
A = zeros(3,16);
A(1,I1) = 1; A(1,HE1) = -1; b(1) = 132000;
A(2,EP) = -1; A(2,PP) = -1; b(2) = -12000;
A(3,[P1,P2,PP]) = [-1,-1,-1];
b(3) = -24550;
```

Write Linear Equality Constraints

There are eight linear equations in the equations of “Model Description” on page 1-13:

$$\begin{aligned} I2 &= LE2 + HE2 \\ LPS &= LE1 + LE2 + BF2 \\ HPS &= I1 + I2 + BF1 \\ HPS &= C + MPS + LPS \\ I1 &= LE1 + HE1 + C \\ MPS &= HE1 + HE2 + BF1 - BF2 \\ 1359.8 I1 &= 1267.8 HE1 + 1251.4 LE1 + 192 C + 3413 P1 \\ 1359.8 I2 &= 1267.8 HE2 + 1251.4 LE2 + 3413 P2. \end{aligned}$$

In order to have the equations in the form $Aeq x = beq$, put all the variables on one side of the equation. The equations become:

```
LE2 + HE2 - I2 = 0
LE1 + LE2 + BF2 - LPS = 0
I1 + I2 + BF1 - HPS = 0
C + MPS + LPS - HPS = 0
LE1 + HE1 + C - I1 = 0
HE1 + HE2 + BF1 - BF2 - MPS = 0
1267.8 HE1 + 1251.4 LE1 + 192 C + 3413 P1 - 1359.8 I1 = 0
1267.8 HE2 + 1251.4 LE2 + 3413 P2 - 1359.8 I2 = 0.
```

Now write the Aeq matrix and beq vector corresponding to these equations. In your MATLAB workspace, create the Aeq matrix as an 8-by-16 zero matrix, corresponding to 8 linear equations in 16 variables. Create the beq vector with eight components, all zero.

```
Aeq = zeros(8,16); beq = zeros(8,1);
Aeq(1,[LE2,HE2,I2]) = [1,1,-1];
Aeq(2,[LE1,LE2,BF2,LPS]) = [1,1,1,-1];
Aeq(3,[I1,I2,BF1,HPS]) = [1,1,1,-1];
Aeq(4,[C,MPS,LPS,HPS]) = [1,1,1,-1];
Aeq(5,[LE1,HE1,C,I1]) = [1,1,1,-1];
Aeq(6,[HE1,HE2,BF1,BF2,MPS]) = [1,1,1,-1,-1];
Aeq(7,[HE1,LE1,C,P1,I1]) = [1267.8,1251.4,192,3413,-1359.8];
Aeq(8,[HE2,LE2,P2,I2]) = [1267.8,1251.4,3413,-1359.8];
```

Write the Objective

The objective function is

$$f^T x = 0.002614 \text{ HPS} + 0.0239 \text{ PP} + 0.009825 \text{ EP}.$$

Write this expression as a vector f of multipliers of the x vector:

```
f = zeros(size(variables));
f([HPS PP EP]) = [0.002614 0.0239 0.009825];
```

Solve the Problem with linprog

You now have inputs required by the linprog solver. Call the solver and print the outputs in formatted form:

```
[x fval] = linprog(f,A,b,Aeq,beq,lb,ub);
```

```

for d = 1:N
    fprintf('%12.2f \t%s\n',x(d),variables{d})
end
fval

```

The result:

Optimization terminated.

```

136328.74 I1
244000.00 I2
128159.00 HE1
143377.00 HE2
    0.00 LE1
100623.00 LE2
    8169.74 C
    0.00 BF1
    0.00 BF2
380328.74 HPS
271536.00 MPS
100623.00 LPS
    6250.00 P1
    7060.71 P2
11239.29 PP
    760.71 EP

```

```

fval =
    1.2703e+003

```

Examine the Solution

The `fval` output gives the smallest value of the objective function at any feasible point.

The solution vector `x` is the point where the objective function has the smallest value. Notice that:

- `BF1`, `BF2`, and `LE1` are 0, their lower bounds.
- `I2` is 244,000, its upper bound.
- The nonzero components of the `f` vector are

- HPS — 380,328.74
- PP — 11,239.29
- EP — 760.71

The video Optimization Modeling 2 gives interpretations of these characteristics in terms of the original problem.

Bibliography

[1] Edgar, Thomas F., and David M. Himmelblau. *Optimization of Chemical Processes*. McGraw-Hill, New York, 1988.

Setting Up an Optimization

- “Optimization Toolbox Solvers” on page 2-2
- “Choosing a Solver” on page 2-4
- “Writing Objective Functions” on page 2-18
- “Matrix Arguments” on page 2-32
- “Writing Constraints” on page 2-33
- “Objective and Nonlinear Constraints in the Same Function” on page 2-48
- “Passing Extra Parameters” on page 2-53
- “What Are Options?” on page 2-57
- “Options in Common Use: Tuning and Troubleshooting” on page 2-58
- “Set and Change Options” on page 2-59
- “Choose Between `optimoptions` and `optimset`” on page 2-60
- “View Options” on page 2-63
- “Tolerances and Stopping Criteria” on page 2-65
- “Checking Validity of Gradients or Jacobians” on page 2-69

Optimization Toolbox Solvers

There are four general categories of Optimization Toolbox solvers:

- Minimizers

This group of solvers attempts to find a local minimum of the objective function near a starting point x_0 . They address problems of unconstrained optimization, linear programming, quadratic programming, and general nonlinear programming.

- Multiobjective minimizers

This group of solvers attempts to either minimize the maximum value of a set of functions (`fminimax`), or to find a location where a collection of functions is below some prespecified values (`fgoalattain`).

- Equation solvers

This group of solvers attempts to find a solution to a scalar- or vector-valued nonlinear equation $f(x) = 0$ near a starting point x_0 . Equation-solving can be considered a form of optimization because it is equivalent to finding the minimum norm of $f(x)$ near x_0 .

- Least-Squares (curve-fitting) solvers

This group of solvers attempts to minimize a sum of squares. This type of problem frequently arises in fitting a model to data. The solvers address problems of finding nonnegative solutions, bounded or linearly constrained solutions, and fitting parameterized nonlinear models to data.

For more information see “Problems Handled by Optimization Toolbox Functions” on page 2-14. See “Optimization Decision Table” on page 2-4 for aid in choosing among solvers for minimization.

Minimizers formulate optimization problems in the form

$$\min_x f(x),$$

possibly subject to constraints. $f(x)$ is called an *objective function*. In general, $f(x)$ is a scalar function of type `double`, and x is a vector or scalar of type `double`. However, multiobjective optimization, equation solving, and some sum-of-squares minimizers, can have vector or matrix objective functions $F(x)$

of type `double`. To use Optimization Toolbox solvers for maximization instead of minimization, see “Maximizing an Objective” on page 2-31.

Write the objective function for a solver in the form of a function file or anonymous function handle. You can supply a gradient $\nabla f(x)$ for many solvers, and you can supply a Hessian for several solvers. See “Writing Objective Functions” on page 2-18. Constraints have a special form, as described in “Writing Constraints” on page 2-33.

Choosing a Solver

In this section...
“Optimization Decision Table” on page 2-4
“Choosing the Algorithm” on page 2-7
“Problems Handled by Optimization Toolbox Functions” on page 2-14

Optimization Decision Table

The following table is designed to help you choose a solver. It does not address multiobjective optimization or equation solving. There are more details on all the solvers in “Problems Handled by Optimization Toolbox Functions” on page 2-14.

Use the table as follows:

1 Identify your objective function as one of five types:

- Linear
- Quadratic
- Sum-of-squares (Least squares)
- Smooth nonlinear
- Nonsmooth

2 Identify your constraints as one of five types:

- None (unconstrained)
- Bound
- Linear (including bound)
- General smooth
- Discrete (integer)

3 Use the table to identify a relevant solver.

In this table:

- * means relevant solvers are found in Global Optimization Toolbox functions (licensed separately from Optimization Toolbox solvers).
- `fmincon` applies to most smooth objective functions with smooth constraints. It is not listed as a preferred solver for least squares or linear or quadratic programming because the listed solvers are usually more efficient.
- The table has suggested functions, but it is not meant to unduly restrict your choices. For example, `fmincon` can be effective on some nonsmooth problems.
- The Global Optimization Toolbox `ga` function can address mixed-integer programming problems.

Solvers by Objective and Constraint

Constraint Type	Objective Type				
	Linear	Quadratic	Least Squares	Smooth nonlinear	Nonsmooth
None	n/a ($f = \text{const}$, or $\min = -\infty$)	quadprog, Information	\, lsqcurvefit, lsqnonlin, Information	fminsearch, fminunc, Information	fminsearch, *
Bound	linprog, Information	quadprog, Information	lsqcurvefit, lsqlin, lsqnonlin, lsqnonneg, Information	fminbnd, fmincon, fseminf, Information	fminbnd, *
Linear	linprog, Information	quadprog, Information	lsqlin, Information	fmincon, fseminf, Information	*
General smooth	fmincon, Information	fmincon, Information	fmincon, Information	fmincon, fseminf, Information	*
Discrete	intlinprog, Information	*	*	*	*

Note This table does not list multiobjective solvers nor equation solvers. See “Problems Handled by Optimization Toolbox Functions” on page 2-14 for a complete list of problems addressed by Optimization Toolbox functions.

Note Some solvers have several algorithms. For help choosing, see “Choosing the Algorithm” on page 2-7.

Choosing the Algorithm

- “fmincon Algorithms” on page 2-7
- “fsolve Algorithms” on page 2-8
- “fminunc Algorithms” on page 2-9
- “Least Squares Algorithms” on page 2-10
- “Linear Programming Algorithms” on page 2-11
- “Quadratic Programming Algorithms” on page 2-11
- “Large-Scale vs. Medium-Scale Algorithms” on page 2-12
- “Potential Inaccuracy with Interior-Point Algorithms” on page 2-12

fmincon Algorithms

fmincon has four algorithm options:

- 'interior-point' (default)
- 'trust-region-reflective'
- 'sqp'
- 'active-set'

Use optimoptions to set the Algorithm option at the command line.

Recommendations

- Use the 'interior-point' algorithm first.
For help if the minimization fails, see “When the Solver Fails” on page 4-3 or “When the Solver Might Have Succeeded” on page 4-15.
- To run an optimization again to obtain more speed on small- to medium-sized problems, try 'sqp' next, and 'active-set' last.
- Use 'trust-region-reflective' when applicable. Your problem must have: objective function includes gradient, only bounds, or only linear equality constraints (but not both).

See “Potential Inaccuracy with Interior-Point Algorithms” on page 2-12.

Reasoning Behind the Recommendations.

- 'interior-point' handles large, sparse problems, as well as small dense problems. The algorithm satisfies bounds at all iterations, and can recover from NaN or Inf results. It is a large-scale algorithm; see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12. The algorithm can use special techniques for large-scale problems. For details, see “Interior-Point Algorithm” on page 10-60.
- 'sqp' satisfies bounds at all iterations. The algorithm can recover from NaN or Inf results. It is not a large-scale algorithm; see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12.
- 'active-set' can take large steps, which adds speed. The algorithm is effective on some problems with nonsmooth constraints. It is not a large-scale algorithm; see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12.
- 'trust-region-reflective' requires you to provide a gradient, and allows only bounds or linear equality constraints, but not both. Within these limitations, the algorithm handles both large sparse problems and small dense problems efficiently. It is a large-scale algorithm; see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12. The algorithm can use special techniques to save memory usage, such as a Hessian multiply function. For details, see “Trust-Region-Reflective Algorithm” on page 10-57.

fsolve Algorithms

fsolve has three algorithms:

- 'trust-region-dogleg' (default)
- 'trust-region-reflective'
- 'levenberg-marquardt'

Use optimoptions to set the Algorithm option at the command line.

Recommendations

- Use the 'trust-region-dogleg' algorithm first.
For help if `fsolve` fails, see “When the Solver Fails” on page 4-3 or “When the Solver Might Have Succeeded” on page 4-15.
- To solve equations again if you have a Jacobian multiply function, or want to tune the internal algorithm (see “Trust-Region-Reflective Algorithm Only” on page 10-134), try 'trust-region-reflective'.
- Try timing all the algorithms, including 'levenberg-marquardt', to find the algorithm that works best on your problem.

Reasoning Behind the Recommendations.

- 'trust-region-dogleg' is the only algorithm that is specially designed to solve nonlinear equations. The others attempt to minimize the sum of squares of the function.
- The 'trust-region-reflective' algorithm is effective on sparse problems. It can use special techniques such as a Jacobian multiply function for large-scale problems.

fminunc Algorithms

`fminunc` has two algorithms:

- 'trust-region' (formerly `LargeScale = 'on'`), the default
- 'quasi-newton' (formerly `LargeScale = 'off'`)

Use `optimoptions` to set the `Algorithm` option at the command line.

Recommendations

- If your objective function includes a gradient, use `'Algorithm' = 'trust-region'`.
- Otherwise, use `'Algorithm' = 'quasi-newton'`.

For help if the minimization fails, see “When the Solver Fails” on page 4-3 or “When the Solver Might Have Succeeded” on page 4-15.

Least Squares Algorithms

lsqin. lsqin has two algorithms:

- 'trust-region-reflective' (formerly LargeScale = 'on'), the default
- 'active-set' (formerly LargeScale = 'off')

Use optimoptions to set the Algorithm option at the command line.

Recommendations

- If you have no constraints or only bound constraints, use 'trust-region-reflective'.
- If you have linear constraints, use 'active-set'.

For help if the minimization fails, see “When the Solver Fails” on page 4-3 or “When the Solver Might Have Succeeded” on page 4-15.

lsqcurvefit and lsqnonlin. lsqcurvefit and lsqnonlin have two algorithms:

- 'trust-region-reflective' (default)
- 'levenberg-marquardt'

Use optimoptions to set the Algorithm option at the command line.

Recommendations

- If you have no constraints or only bound constraints, use 'trust-region-reflective'.
- If your problem is underdetermined (fewer equations than dimensions), use 'levenberg-marquardt'.

For help if the minimization fails, see “When the Solver Fails” on page 4-3 or “When the Solver Might Have Succeeded” on page 4-15.

Linear Programming Algorithms

linprog has three algorithms:

- 'interior-point' (formerly LargeScale = 'on'), the default
- 'active-set' (formerly LargeScale = 'off')
- 'simplex' (formerly LargeScale = 'off', Simplex = 'on')

Use optimoptions to set the Algorithm option at the command line.

Recommendations

Use the interior-point algorithm.

For help if the minimization fails, see “When the Solver Fails” on page 4-3 or “When the Solver Might Have Succeeded” on page 4-15.

See “Potential Inaccuracy with Interior-Point Algorithms” on page 2-12.

Quadratic Programming Algorithms

quadprog has three algorithms:

- 'interior-point-convex' (default)
- 'trust-region-reflective' (formerly LargeScale = 'on')
- 'active-set' (formerly LargeScale = 'off').

Use optimoptions to set the Algorithm option at the command line.

Recommendations

- If you have a convex problem, or if you don't know whether your problem is convex, use 'interior-point-convex'.
- If you have a nonconvex problem with only bounds, or with only linear equalities, use 'trust-region-reflective'.
- If you have a nonconvex problem that does not satisfy the restrictions of 'trust-region-reflective', use 'active-set'.

Recommendations

For help if the minimization fails, see “When the Solver Fails” on page 4-3 or “When the Solver Might Have Succeeded” on page 4-15.

See “Potential Inaccuracy with Interior-Point Algorithms” on page 2-12.

Large-Scale vs. Medium-Scale Algorithms

An optimization algorithm is *large scale* when it uses linear algebra that does not need to store, nor operate on, full matrices. This may be done internally by storing sparse matrices, and by using sparse linear algebra for computations whenever possible. Furthermore, the internal algorithms either preserve sparsity, such as a sparse Cholesky decomposition, or do not generate matrices, such as a conjugate gradient method.

In contrast, *medium-scale* methods internally create full matrices and use dense linear algebra. If a problem is sufficiently large, full matrices take up a significant amount of memory, and the dense linear algebra may require a long time to execute.

Don't let the name “large scale” mislead you; you can use a large-scale algorithm on a small problem. Furthermore, you do not need to specify any sparse matrices to use a large-scale algorithm. Choose a medium-scale algorithm to access extra functionality, such as additional constraint types, or possibly for better performance.

Potential Inaccuracy with Interior-Point Algorithms

Interior-point algorithms in `fmincon`, `quadprog`, and `linprog` have many good characteristics, such as low memory usage and the ability to solve large problems quickly. However, their solutions can be slightly less accurate than those from other algorithms. The reason for this potential inaccuracy is that the (internally calculated) barrier function keeps iterates away from inequality constraint boundaries.

For most practical purposes, this inaccuracy is usually quite small.

To reduce the inaccuracy, try to:

- Rerun the solver with smaller TolX, TolFun, and possibly TolCon tolerances (but keep the tolerances sensible.) See “Tolerances and Stopping Criteria” on page 2-65).
- Run a different algorithm, starting from the interior-point solution. This can fail, because some algorithms can use excessive memory or time, and some linprog and quadprog algorithms do not accept an initial point.

For example, try to minimize the function x when bounded below by 0. Using the fmincon interior-point algorithm:

```
options = optimoptions(@fmincon,'Algorithm','interior-point','Display','off');
x = fmincon(@(x)x,1,[],[],[],[],0,[],[],options)
```

```
x =
```

```
2.0000e-08
```

Using the fmincon sqp algorithm:

```
options.Algorithm = 'sqp';
x2 = fmincon(@(x)x,1,[],[],[],[],0,[],[],options)
```

```
x2 =
```

```
0
```

Similarly, solve the same problem using the linprog interior-point algorithm:

```
opts = optimoptions(@linprog,'Display','off','Algorithm','interior-point');
x = linprog(1,[],[],[],[],0,[],1,opts)
```

```
x =
```

```
2.0833e-13
```

Using the linprog simplex algorithm:

```
opts.Algorithm = 'simplex';
x2 = linprog(1,[],[],[],[],0,[],1,opts)
```

```
x2 =
```

0

In these cases, the interior-point algorithms are less accurate, but the answers are quite close to the correct answer.

Problems Handled by Optimization Toolbox Functions

The following tables show the functions available for minimization, equation solving, multiobjective optimization, and solving least-squares or data-fitting problems.

Minimization Problems

Type	Formulation	Solver
Scalar minimization	$\min_x f(x)$ such that $lb < x < ub$ (x is scalar)	fminbnd
Unconstrained minimization	$\min_x f(x)$	fminunc, fminsearch
Linear programming	$\min_x f^T x$ such that $Ax \leq b, Aeqx = beq, lb \leq x \leq ub$	linprog
Mixed-integer linear programming	$\min_x f^T x$ such that $Ax \leq b, Aeqx = beq, lb \leq x \leq ub,$ $x(\text{intcon})$ is integer-valued.	intlinprog
Quadratic programming	$\min_x \frac{1}{2} x^T Hx + c^T x$ such that $Ax \leq b, Aeqx = beq, lb \leq x \leq ub$	quadprog

Minimization Problems (Continued)

Type	Formulation	Solver
Constrained minimization	$\min_x f(x)$ <p>such that $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, $lb \leq x \leq ub$</p>	fmincon
Semi-infinite minimization	$\min_x f(x)$ <p>such that $K(x,w) \leq 0$ for all w, $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, $lb \leq x \leq ub$</p>	fseminf

Multiobjective Problems

Type	Formulation	Solver
Goal attainment	$\min_{x,\gamma} \gamma$ <p>such that $F(x) - w\gamma \leq \text{goal}$, $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, $lb \leq x \leq ub$</p>	fgoalattain
Minimax	$\min_x \max_i F_i(x)$ <p>such that $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, $lb \leq x \leq ub$</p>	fminimax

Equation Solving Problems

Type	Formulation	Solver
Linear equations	$Cx = d$, n equations, n variables	\ (matrix left division)
Nonlinear equation of one variable	$f(x) = 0$	fzero
Nonlinear equations	$F(x) = 0$, n equations, n variables	fsolve

Least-Squares (Model-Fitting) Problems

Type	Formulation	Solver
Linear least-squares	$\min_x \frac{1}{2} \ C \cdot x - d\ _2^2$ m equations, n variables	\ (matrix left division)
Nonnegative linear-least-squares	$\min_x \frac{1}{2} \ C \cdot x - d\ _2^2$ such that $x \geq 0$	lsqnonneg
Constrained linear-least-squares	$\min_x \frac{1}{2} \ C \cdot x - d\ _2^2$ such that $Ax \leq b$, $Aeq\ x = beq$, $lb \leq x \leq ub$	lsqlin

Least-Squares (Model-Fitting) Problems (Continued)

Type	Formulation	Solver
Nonlinear least-squares	$\min_x \ F(x)\ _2^2 = \min_x \sum_i F_i^2(x)$ <p>such that $lb \leq x \leq ub$</p>	lsqnonlin
Nonlinear curve fitting	$\min_x \ F(x, xdata) - ydata\ _2^2$ <p>such that $lb \leq x \leq ub$</p>	lsqcurvefit

Writing Objective Functions

In this section...
“Types of Objective Functions” on page 2-18
“Writing Scalar Objective Functions” on page 2-19
“Writing Vector and Matrix Objective Functions” on page 2-27
“Writing Objective Functions for Linear or Quadratic Problems” on page 2-30
“Maximizing an Objective” on page 2-31

Types of Objective Functions

Many Optimization Toolbox solvers minimize a scalar function of a multidimensional vector. The *objective* function is the function the solvers attempt to minimize. Several solvers accept vector-valued objective functions, and some solvers use objective functions you specify by vectors or matrices.

Objective Type	Solvers	How to Write Objectives
Scalar	fmincon fminunc fminbnd fminsearch fseminf fzero	“Writing Scalar Objective Functions” on page 2-19
Nonlinear least squares	lsqcurvefit lsqnonlin	
Multivariable equation solving	fsolve	
Multiobjective	fgoalattain fminimax	

Objective Type	Solvers	How to Write Objectives
Linear programming	linprog	“Writing Objective Functions for Linear or Quadratic Problems” on page 2-30
Binary integer programming	bintprog	
Linear least squares	lsqlin	
	lsqnonneg	
Quadratic programming	quadprog	

Writing Scalar Objective Functions

- “Function Files” on page 2-19
- “Anonymous Function Objectives” on page 2-21
- “Including Derivatives” on page 2-22

Function Files

A scalar objective function file accepts one input, say x , and returns one scalar output, say f . The input x can be a scalar, vector, or matrix. A function file can return more outputs (see “Including Derivatives” on page 2-22).

For example, suppose your objective is a function of three variables, x , y , and z :

$$f(x) = 3*(x - y)^4 + 4*(x + z)^2 / (1 + x^2 + y^2 + z^2) + \cosh(x - 1) + \tanh(y + z).$$

- 1 Write this function as a file that accepts the vector $x_{in} = [x;y;z]$ and returns f :

```
function f = myObjective(xin)
f = 3*(xin(1)-xin(2))^4 + 4*(xin(1)+xin(3))^2/(1+norm(xin)^2) ...
    + cosh(xin(1)-1) + tanh(xin(2)+xin(3));
```

- 2 Save it as a file named `myObjective.m` to a folder on your MATLAB path.
- 3 Check that the function evaluates correctly:

```
myObjective([1;2;3])
```

```
ans =  
    9.2666
```

For information on how to include extra parameters, see “Passing Extra Parameters” on page 2-53. For more complex examples of function files, see “Minimization with Gradient and Hessian Sparsity Pattern” on page 6-22 or “Minimization with Bound Constraints and Banded Preconditioner” on page 6-72.

Local Functions and Nested Functions. Functions can exist inside other files as local functions or nested functions. Using local functions or nested functions can lower the number of distinct files you save. Using nested functions also lets you access extra parameters, as shown in “Nested Functions” on page 2-55.

For example, suppose you want to minimize the `myObjective.m` objective function, described in “Function Files” on page 2-19, subject to the `ellipseparabola.m` constraint, described in “Nonlinear Constraints” on page 2-37. Instead of writing two files, `myObjective.m` and `ellipseparabola.m`, write one file that contains both functions as local functions:

```
function [x fval] = callObjConstr(x0,options)  
% Using a local function for just one file  
  
if nargin < 2  
    options = optimoptions('fmincon','Algorithm','interior-point');  
end  
  
[x fval] = fmincon(@myObjective,x0,[],[],[],[],[],[], ...  
    @ellipseparabola,options);  
  
function f = myObjective(xin)  
f = 3*(xin(1)-xin(2))^4 + 4*(xin(1)+xin(3))^2/(1+sum(xin.^2)) ...  
    + cosh(xin(1)-1) + tanh(xin(2)+xin(3));  
  
function [c,ceq] = ellipseparabola(x)  
c(1) = (x(1)^2)/9 + (x(2)^2)/4 - 1;  
c(2) = x(1)^2 - x(2) - 1;  
ceq = [];
```

Solve the constrained minimization starting from the point [1;1;1]:

```
[x fval] = callObjConstr(ones(3,1))
```

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.

```
x =
    1.1835
    0.8345
   -1.6439
```

```
fval =
    0.5383
```

Anonymous Function Objectives

Use anonymous functions to write simple objective functions. For more information about anonymous functions, see “What Are Anonymous Functions?” in the MATLAB Programming Fundamentals documentation. Rosenbrock’s function is simple enough to write as an anonymous function:

```
anonrosen = @(x)(100*(x(2) - x(1)^2)^2 + (1-x(1))^2);
```

Check that anonrosen evaluates correctly at [-1 2]:

```
anonrosen([-1 2])
```

```
ans =
    104
```

Minimizing anonrosen with fminunc yields the following results:

```
options = optimoptions(@fminunc,'Algorithm','quasi-newton');
[x fval] = fminunc(anonrosen,[-1;2],options)
```

Local minimum found.

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```
x =
    1.0000
    1.0000

fval =
    1.2262e-10
```

Including Derivatives

For `fmincon` and `fminunc`, you can include gradients in the objective function. You can also include Hessians, depending on the algorithm. The Hessian matrix $H_{i,j}(x) = \partial^2 f / \partial x_i \partial x_j$.

The following table shows which algorithms can use gradients and Hessians.

Solver	Algorithm	Gradient	Hessian
fmincon	active-set	Optional	No
	interior-point	Optional	Optional separate function (see “Hessian” on page 10-49)
	sqp	Optional	No
	trust-region-reflective	Required	Optional
fminunc	trust-region	Required	Optional
	quasi-newton	Optional	No

- “Benefits of Including Derivatives” on page 2-23
- “Choose Input Hessian for interior-point fmincon” on page 2-23
- “How to Include Derivatives” on page 2-24

Benefits of Including Derivatives. If you do not provide gradients, solvers estimate gradients via finite differences. If you provide gradients, your solver need not perform this finite difference estimation, so can save time and be more accurate. Furthermore, solvers use an approximate Hessian, which can be far from the true Hessian. Providing a Hessian can yield a solution in fewer iterations.

For constrained problems, providing a gradient has another advantage. A solver can reach a point x such that x is feasible, but, for this x , finite differences around x always lead to an infeasible point. Suppose further that the objective function at an infeasible point returns a complex output, Inf, NaN, or error. In this case, a solver can fail or halt prematurely. Providing a gradient allows a solver to proceed. To obtain this benefit, you might also need to include the gradient of a nonlinear constraint function, and set the GradConstr option to 'on'. See “Nonlinear Constraints” on page 2-37.

Choose Input Hessian for interior-point fmincon. The fmincon interior-point algorithm has many options for selecting an input Hessian. For syntax details, see “Hessian” on page 10-49. Here are the options, along with estimates of their relative characteristics.

Hessian	Relative Memory Usage	Relative Efficiency
'bfgs' (default)	High (for large problems)	High
'lbfgs'	Low to Moderate	Moderate
'fin-diff-grads'	Low	Moderate
'user-supplied' with 'HessMult'	Low (can depend on your code)	Moderate
'user-supplied' with 'HessFcn'	? (depends on your code)	High (depends on your code)

Use the default 'bfgs' Hessian unless you

- Run out of memory — Try 'lbfgs' instead of 'bfgs'. If you can provide your own gradients, try 'fin-diff-grads', and set the GradObj and GradConstr options to 'on'.

- Want more efficiency — Provide your own gradients and Hessian. See “fmincon Interior-Point Algorithm with Analytic Hessian” on page 6-57 and “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85.

The reason 'lbfgs' has only moderate efficiency is twofold. It has relatively expensive Sherman-Morrison updates. And the resulting iteration step can be somewhat inaccurate due to the 'lbfgs' limited memory.

The reason 'fin-diff-grads' and HessMult have only moderate efficiency is that they use a conjugate gradient approach. They accurately estimate the Hessian of the objective function, but they do not generate the most accurate iteration step. For more information, see “fmincon Interior Point Algorithm” on page 6-43, and its discussion of the LDL approach and the conjugate gradient approach to solving Equation 6-53.

How to Include Derivatives.

- 1 Write code that returns:
 - The objective function (scalar) as the first output
 - The gradient (vector) as the second output
 - Optionally, the Hessian (matrix) as the third output
- 2 Set the GradObj option to 'on' with optimoptions.
- 3 Optionally, set the Hessian option to 'on' or 'user-supplied'.

For the fmincon interior-point solver, set the Hessian option to 'user-supplied' and set the 'HessFcn' option to `@hessianfcn`, where `hessianfcn` is a function that computes the Hessian of the Lagrangian. For details, see “Hessian” on page 10-49. For an example, see “fmincon Interior-Point Algorithm with Analytic Hessian” on page 6-57.

- 4 Optionally, check if your gradient function matches a finite-difference approximation. See “Checking Validity of Gradients or Jacobians” on page 2-69.

Tip For most flexibility, write *conditionalized* code. Conditionalized means that the number of function outputs can vary, as shown in the following example. Conditionalized code does not error depending on the value of the `GradObj` or `Hessian` option. Unconditionalized code requires you to set these options appropriately.

For example, consider Rosenbrock's function

$$f(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2,$$

which is described and plotted in “Solve a Constrained Nonlinear Problem” on page 1-3. The gradient of $f(x)$ is

$$\nabla f(x) = \begin{bmatrix} -400(x_2 - x_1^2)x_1 - 2(1 - x_1) \\ 200(x_2 - x_1^2) \end{bmatrix},$$

and the Hessian $H(x)$ is

$$H(x) = \begin{bmatrix} 1200x_1^2 - 400x_2 + 2 & -400x_1 \\ -400x_1 & 200 \end{bmatrix}.$$

`rosenthree` is an unconditionalized function that returns the Rosenbrock function with its gradient and Hessian:

```
function [f g H] = rosenthree(x)
% Calculate objective f, gradient g, Hessian H
f = 100*(x(2) - x(1)^2)^2 + (1-x(1))^2;
g = [-400*(x(2) - x(1)^2)*x(1) - 2*(1-x(1));
     200*(x(2) - x(1)^2)];
H = [1200*x(1)^2 - 400*x(2) + 2, -400*x(1);
     -400*x(1), 200];
```

`rosenboth` is a conditionalized function that returns whatever the solver requires:

```
function [f g H] = rosenboth(x)
% Calculate objective f
f = 100*(x(2) - x(1)^2)^2 + (1-x(1))^2;

if nargout > 1 % gradient required
    g = [-400*(x(2)-x(1)^2)*x(1)-2*(1-x(1));
        200*(x(2)-x(1)^2)];

    if nargout > 2 % Hessian required
        H = [1200*x(1)^2-400*x(2)+2, -400*x(1);
            -400*x(1), 200];
    end
end

end
```

nargout checks the number of arguments that a calling function specifies. See “Find Number of Function Arguments” in the MATLAB Programming Fundamentals documentation.

The `fminunc` solver, designed for unconstrained optimization, allows you to minimize Rosenbrock’s function. Tell `fminunc` to use the gradient and Hessian by setting options:

```
options = optimoptions(@fminunc,'GradObj','on','Hessian','on');
```

Run `fminunc` starting at `[-1;2]`:

```
[x fval] = fminunc(@rosenboth,[-1;2],options)
Local minimum found.
```

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```
x =
    1.0000
    1.0000

fval =
    1.9310e-017
```

If you have a Symbolic Math Toolbox™ license, you can calculate gradients and Hessians automatically, as described in “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85.

Writing Vector and Matrix Objective Functions

Some solvers, such as `fsolve` and `lsqcurvefit`, have objective functions that are vectors or matrices. The main difference in usage between these types of objective functions and scalar objective functions is the way to write their derivatives. The first-order partial derivatives of a vector-valued or matrix-valued function is called a Jacobian; the first-order partial derivatives of a scalar function is called a gradient.

- “Jacobians of Vector Functions” on page 2-27
- “Jacobians of Matrix Functions” on page 2-28
- “Jacobians with Matrix-Valued Independent Variables” on page 2-29

Jacobians of Vector Functions

If x is a vector of independent variables, and $F(x)$ is a vector function, the Jacobian $J(x)$ is

$$J_{ij}(x) = \frac{\partial F_i(x)}{\partial x_j}.$$

If F has m components, and x has k components, J is an m -by- k matrix.

For example, if

$$F(x) = \begin{bmatrix} x_1^2 + x_2 x_3 \\ \sin(x_1 + 2x_2 - 3x_3) \end{bmatrix},$$

then $J(x)$ is

$$J(x) = \begin{bmatrix} 2x_1 & x_3 & x_2 \\ \cos(x_1 + 2x_2 - 3x_3) & 2\cos(x_1 + 2x_2 - 3x_3) & -3\cos(x_1 + 2x_2 - 3x_3) \end{bmatrix}.$$

The function file associated with this example is:

```

function [F jacF] = vectorObjective(x)
F = [x(1)^2 + x(2)*x(3);
     sin(x(1) + 2*x(2) - 3*x(3))];
if nargin > 1 % need Jacobian
    jacF = [2*x(1), x(3), x(2);
           cos(x(1)+2*x(2)-3*x(3)), 2*cos(x(1)+2*x(2)-3*x(3)), ...
           -3*cos(x(1)+2*x(2)-3*x(3))];
end

```

Jacobians of Matrix Functions

The Jacobian of a matrix $F(x)$ is defined by changing the matrix to a vector, column by column. For example, rewrite the matrix

$$F = \begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \\ F_{31} & F_{32} \end{bmatrix}$$

as a vector f :

$$f = \begin{bmatrix} F_{11} \\ F_{21} \\ F_{31} \\ F_{12} \\ F_{22} \\ F_{32} \end{bmatrix}.$$

The Jacobian of F is as the Jacobian of f ,

$$J_{ij} = \frac{\partial f_i}{\partial x_j}.$$

If F is an m -by- n matrix, and x is a k -vector, the Jacobian is an mn -by- k matrix.

For example, if

$$F(x) = \begin{bmatrix} x_1 x_2 & x_1^3 + 3x_2^2 \\ 5x_2 - x_1^4 & x_2 / x_1 \\ 4 - x_2^2 & x_1^3 - x_2^4 \end{bmatrix},$$

then the Jacobian of F is

$$J(x) = \begin{bmatrix} x_2 & x_1 \\ -4x_1^3 & 5 \\ 0 & -2x_2 \\ 3x_1^2 & 6x_2 \\ -x_2 / x_1^2 & 1/x_1 \\ 3x_1^2 & -4x_2^3 \end{bmatrix}.$$

Jacobians with Matrix-Valued Independent Variables

If x is a matrix, define the Jacobian of $F(x)$ by changing the matrix x to a vector, column by column. For example, if

$$X = \begin{bmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{bmatrix},$$

then the gradient is defined in terms of the vector

$$x = \begin{bmatrix} x_{11} \\ x_{21} \\ x_{12} \\ x_{22} \end{bmatrix}.$$

With

$$F = \begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \\ F_{31} & F_{32} \end{bmatrix},$$

and with f the vector form of F as above, the Jacobian of $F(X)$ is defined as the Jacobian of $f(x)$:

$$J_{ij} = \frac{\partial f_i}{\partial x_j}.$$

So, for example,

$$J(3,2) = \frac{\partial f(3)}{\partial x(2)} = \frac{\partial F_{31}}{\partial X_{21}}, \text{ and } J(5,4) = \frac{\partial f(5)}{\partial x(4)} = \frac{\partial F_{22}}{\partial X_{22}}.$$

If F is an m -by- n matrix and x is a j -by- k matrix, then the Jacobian is an mn -by- jk matrix.

Writing Objective Functions for Linear or Quadratic Problems

The following solvers handle linear or quadratic objective functions:

- `linprog` and `intlinprog`: minimize

$$f'x = f(1)*x(1) + f(2)*x(2) + \dots + f(n)*x(n).$$

Input the vector f for the objective. See the examples in “Linear Programming and Mixed-Integer Linear Programming”.

- `lsqlin` and `lsqnonneg`: minimize

$$\|Cx - d\|.$$

Input the matrix C and the vector d for the objective. See “Linear Least Squares with Bound Constraints” on page 6-219.

- `quadprog`: minimize

$$\begin{aligned} & 1/2 * x'Hx + f'x \\ & = 1/2 * (x(1)*H(1,1)*x(1) + 2*x(1)*H(1,2)*x(2) + \dots \\ & + x(n)*H(n,n)*x(n)) + f(1)*x(1) + f(2)*x(2) + \dots + f(n)*x(n). \end{aligned}$$

Input both the vector f and the symmetric matrix H for the objective. See “Quadratic Programming”.

Maximizing an Objective

All solvers attempt to minimize an objective function. If you have a maximization problem, that is, a problem of the form

$$\max_x f(x),$$

then define $g(x) = -f(x)$, and minimize g .

For example, to find the maximum of $\tan(\cos(x))$ near $x = 5$, evaluate:

```
[x fval] = fminunc(@(x) -tan(cos(x)),5)
Local minimum found.
```

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```
x =
    6.2832
```

```
fval =
   -1.5574
```

The maximum is 1.5574 (the negative of the reported fval), and occurs at $x = 6.2832$. This answer is correct since, to five digits, the maximum is $\tan(1) = 1.5574$, which occurs at $x = 2\pi = 6.2832$.

Matrix Arguments

Solvers accept matrix initial point `x0`, where matrix means an array of any size. They also accept matrix bounds `lb` and `ub`. Here's how solvers handle matrix arguments.

- Internally, solvers convert matrix arguments into vectors before processing. For example, `x0` becomes `x0(:)`. For an explanation of this syntax, see the `A(:)` entry in `colon`.
- For output, solvers reshape the solution `x` to the same size as the input `x0`.
- When `x0` is a matrix, solvers pass `x` as a matrix of the same size as `x0` to both the objective function and to any nonlinear constraint function.
- Linear constraints, though, take `x` in vector form, `x(:)`. In other words, a linear constraint of the form

$$A*x \leq b \text{ or } Aeq*x = beq$$

takes `x` as a vector, not a matrix. Ensure that your matrix `A` or `Aeq` has the same number of columns as `x0` has elements, or the solver will error.

Writing Constraints

In this section...

- “Types of Constraints” on page 2-33
- “Iterations Can Violate Constraints” on page 2-34
- “Bound Constraints” on page 2-35
- “Linear Inequality Constraints” on page 2-36
- “Linear Equality Constraints” on page 2-37
- “Nonlinear Constraints” on page 2-37
- “Or Instead of And Constraints” on page 2-41
- “How to Use All Types of Constraints” on page 2-46

Types of Constraints

Optimization Toolbox solvers have special forms for constraints:

- “Bound Constraints” on page 2-35 — Lower and upper bounds on individual components: $x \geq l$ and $x \leq u$.
- “Linear Inequality Constraints” on page 2-36 — $Ax \leq b$. A is an m -by- n matrix, which represents m constraints for an n -dimensional vector x . b is m -dimensional.
- “Linear Equality Constraints” on page 2-37 — $Aeqx = beq$. Equality constraints have the same form as inequality constraints.
- “Nonlinear Constraints” on page 2-37 — $c(x) \leq 0$ and $ceq(x) = 0$. Both c and ceq are scalars or vectors representing several constraints.

Optimization Toolbox functions assume that inequality constraints are of the form $c_i(x) \leq 0$ or $Ax \leq b$. Express greater-than constraints as less-than constraints by multiplying them by -1 . For example, a constraint of the form $c_i(x) \geq 0$ is equivalent to the constraint $-c_i(x) \leq 0$. A constraint of the form $Ax \geq b$ is equivalent to the constraint $-Ax \leq -b$. For more information, see “Linear Inequality Constraints” on page 2-36 and “Nonlinear Constraints” on page 2-37.

You can sometimes write constraints in several ways. For best results, use the lowest numbered constraints possible:

- 1 Bounds
- 2 Linear equalities
- 3 Linear inequalities
- 4 Nonlinear equalities
- 5 Nonlinear inequalities

For example, with a constraint $5x \leq 20$, use a bound $x \leq 4$ instead of a linear inequality or nonlinear inequality.

For information on how to pass extra parameters to constraint functions, see “Passing Extra Parameters” on page 2-53.

Iterations Can Violate Constraints

Be careful when writing your objective and constraint functions. Intermediate iterations can lead to points that are infeasible (do not satisfy constraints). If you write objective or constraint functions that assume feasibility, these functions can error or give unexpected results.

For example, if you take a square root or logarithm of x , and $x < 0$, the result is not real. You can try to avoid this error by setting 0 as a lower bound on x . Nevertheless, an intermediate iteration can violate this bound.

Algorithms That Satisfy Bound Constraints

Some solver algorithms satisfy bound constraints at every iteration:

- `fmincon` interior-point, sqp, and trust-region-reflective algorithms
- `lsqcurvefit` trust-region-reflective algorithm
- `lsqnonlin` trust-region-reflective algorithm
- `fminbnd`

Note If you set a lower bound equal to an upper bound, iterations can violate these constraints.

Solvers and Algorithms That Can Violate Bound Constraints

The following solvers and algorithms can violate bound constraints at intermediate iterations:

- `fmincon` active-set algorithm
- `fgoalattain` solver
- `fminimax` solver
- `fseminf` solver

Bound Constraints

Lower and upper bounds limit the components of the solution x .

If you know bounds on the location of an optimum, you can obtain faster and more reliable solutions by explicitly including these bounds in your problem formulation.

Give bounds as vectors with the same length as x , or as matrices with the same number of elements as x .

- If a particular component has no lower bound, use `-Inf` as the bound; similarly, use `Inf` if a component has no upper bound.
- If you have only bounds of one type (upper or lower), you do not need to write the other type. For example, if you have no upper bounds, you do not need to supply a vector of `Infs`.
- If only the first m out of n components have bounds, then you need only supply a vector of length m containing bounds. However, this shortcut causes solvers to throw a warning.

For example, suppose your bounds are:

$$\begin{aligned}x_3 &\geq 8 \\x_2 &\leq 3.\end{aligned}$$

Write the constraint vectors as

$$\begin{aligned}l &= [-\text{Inf}; -\text{Inf}; 8] \\u &= [\text{Inf}; 3] \text{ (throws a warning) or } u = [\text{Inf}; 3; \text{Inf}].\end{aligned}$$

Tip Use `Inf` or `-Inf` instead of a large, arbitrary bound to lower memory usage and increase solver speed. See “Use Inf Instead of a Large, Arbitrary Bound” on page 4-12.

You need not give gradients for bound constraints; solvers calculate them automatically. Bounds do not affect Hessians.

For a more complex example of bounds, see “Set Up a Linear Program” on page 1-13.

Linear Inequality Constraints

Linear inequality constraints have the form $Ax \leq b$. When A is m -by- n , there are m constraints on a variable x with n components. You supply the m -by- n matrix A and the m -component vector b .

Even if you pass an initial point x_0 as a matrix, solvers pass the current point x as a column vector to linear constraints. See “Matrix Arguments” on page 2-32.

For example, suppose that you have the following linear inequalities as constraints:

$$\begin{aligned}x_1 + x_3 &\leq 4, \\2x_2 - x_3 &\geq -2, \\x_1 - x_2 + x_3 - x_4 &\geq 9.\end{aligned}$$

Here $m = 3$ and $n = 4$.

Write these using the following matrix A and vector b :

$$A = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & -2 & 1 & 0 \\ -1 & 1 & -1 & 1 \end{bmatrix},$$

$$b = \begin{bmatrix} 4 \\ 2 \\ -9 \end{bmatrix}.$$

Notice that the “greater than” inequalities were first multiplied by -1 in order to get them into “less than” inequality form. In MATLAB syntax:

```
A = [1 0 1 0;
     0 -2 1 0;
     -1 1 -1 1];
b = [4;2;-9];
```

You do not need to give gradients for linear constraints; solvers calculate them automatically. Linear constraints do not affect Hessians.

For a more complex example of linear constraints, see “Set Up a Linear Program” on page 1-13.

Linear Equality Constraints

Linear equalities have the form $Aeq\ x = beq$, which represents m equations with n -component vector x . You supply the m -by- n matrix Aeq and the m -component vector beq .

You do not need to give gradients for linear constraints; solvers calculate them automatically. Linear constraints do not affect Hessians. The form of this type of constraint is the same as for “Linear Inequality Constraints” on page 2-36.

Nonlinear Constraints

Nonlinear inequality constraints have the form $c(x) \leq 0$, where c is a vector of constraints, one component for each constraint. Similarly, nonlinear equality constraints are of the form $ceq(x) = 0$.

Note Nonlinear constraint functions must return both `c` and `ceq`, the inequality and equality constraint functions, even if they do not both exist. Return empty `[]` for a nonexistent constraint.

For example, suppose that you have the following inequalities as constraints:

$$\frac{x_1^2}{9} + \frac{x_2^2}{4} \leq 1,$$
$$x_2 \geq x_1^2 - 1.$$

Write these constraints in a function file as follows:

```
function [c,ceq]=ellipseparabola(x)
c(1) = (x(1)^2)/9 + (x(2)^2)/4 - 1;
c(2) = x(1)^2 - x(2) - 1;
ceq = [];
end
```

`ellipseparabola` returns empty `[]` for `ceq`, the nonlinear equality function. Also, both inequalities were put into ≤ 0 form.

Including Gradients in Constraint Functions

If you provide gradients for `c` and `ceq`, your solver can run faster and give more reliable results.

Providing a gradient has another advantage. A solver can reach a point `x` such that `x` is feasible, but finite differences around `x` always lead to an infeasible point. In this case, a solver can fail or halt prematurely. Providing a gradient allows a solver to proceed.

To include gradient information, write a conditionalized function as follows:

```
function [c,ceq,gradc,gradceq]=ellipseparabola(x)
c(1) = x(1)^2/9 + x(2)^2/4 - 1;
c(2) = x(1)^2 - x(2) - 1;
ceq = [];

if nargin > 2
```

```

gradc = [2*x(1)/9, 2*x(1); ...
         x(2)/2, -1];
gradceq = [];
end

```

See “Writing Scalar Objective Functions” on page 2-19 for information on conditionalized functions. The gradient matrix has the form

$$\text{gradc}_{i,j} = [\partial c(j)/\partial x_i].$$

The first column of the gradient matrix is associated with $c(1)$, and the second column is associated with $c(2)$. This is the transpose of the form of Jacobians.

To have a solver use gradients of nonlinear constraints, indicate that they exist by using `optimoptions`:

```
options=optimoptions(@fmincon,'GradConstr','on');
```

Make sure to pass the options structure to your solver:

```
[x,fval] = fmincon(@myobj,x0,A,b,Aeq,beq,lb,ub, ...
                 @ellipseparabola,options)
```

If you have a Symbolic Math Toolbox license, you can calculate gradients and Hessians automatically, as described in “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85.

Anonymous Nonlinear Constraint Functions

For information on anonymous objective functions, see “Anonymous Function Objectives” on page 2-21.

Nonlinear constraint functions must return two outputs. The first output corresponds to nonlinear inequalities, and the second corresponds to nonlinear equalities.

Anonymous functions return just one output. So how can you write an anonymous function as a nonlinear constraint?

The `deal` function distributes multiple outputs. For example, suppose your nonlinear inequalities are

$$\frac{x_1^2}{9} + \frac{x_2^2}{4} \leq 1,$$
$$x_2 \geq x_1^2 - 1.$$

Suppose that your nonlinear equality is

$$x_2 = \tanh(x_1).$$

Write a nonlinear constraint function as follows:

```
c = @(x)[x(1)^2/9 + x(2)^2/4 - 1;  
        x(1)^2 - x(2) - 1];  
ceq = @(x)tanh(x(1)) - x(2);  
nonlinfcn = @(x)deal(c(x),ceq(x));
```

To minimize the function $\cosh(x_1) + \sinh(x_2)$ subject to the constraints in `nonlinfcn`, use `fmincon`:

```
obj = @(x)cosh(x(1))+sinh(x(2));  
opts = optimoptions(@fmincon,'Algorithm','sqp');  
z = fmincon(obj,[0;0],[],[],[],[],[],[],[],nonlinfcn,opts)
```

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.

```
z =  
   -0.6530  
   -0.5737
```

To check how well the resulting point `z` satisfies the constraints, use `nonlinfcn`:

```
[cout,ceqout] = nonlinfcn(z)  
  
cout =  
   -0.8704
```


0

```
ceqout =
    1.1102e-016
```

z indeed satisfies all the constraints to within the default value of the TolCon constraint tolerance, $1e-6$.

Or Instead of And Constraints

In general, solvers takes constraints with an implicit AND:

constraint 1 AND constraint 2 AND constraint 3 are all satisfied.

However, sometimes you want an OR:

constraint 1 OR constraint 2 OR constraint 3 is satisfied.

These formulations are not logically equivalent, and there is generally no way to express OR constraints in terms of AND constraints.

Tip Fortunately, nonlinear constraints are extremely flexible. You get OR constraints simply by setting the nonlinear constraint function to the minimum of the constraint functions.

The reason that you can set the minimum as the constraint is due to the nature of “Nonlinear Constraints” on page 2-37: you give them as a set of functions that must be negative at a feasible point. If your constraints are

$$F_1(x) \leq 0 \text{ OR } F_2(x) \leq 0 \text{ OR } F_3(x) \leq 0,$$


then set the nonlinear inequality constraint function $c(x)$ as:

$$c(x) = \min(F_1(x), F_2(x), F_3(x)).$$

$c(x)$ is not smooth, which is a general requirement for constraint functions, due to the minimum. Nevertheless, the method often works.

Note You cannot use the usual bounds and linear constraints in an OR constraint. Instead, convert your bounds and linear constraints to nonlinear constraint functions, as in this example.

For example, suppose your feasible region is the L-shaped region: x is in the rectangle $-1 \leq x(1) \leq 1$, $0 \leq x(2) \leq 1$ OR x is in the rectangle $0 \leq x(1) \leq 1$, $-1 \leq x(2) \leq 1$.



Code for creating the figure

```
% Write the x and y coordinates of the figure, clockwise from (0,0)
```

```

x = [0, -1, -1, 1, 1, 0, 0];
y = [0, 0, 1, 1, -1, -1, 0];
plot(x,y)
xlim([-1.2 1.2])
ylim([-1.2 1.2])
axis equal

```

To represent a rectangle as a nonlinear constraint, instead of as bound constraints, construct a function that is negative inside the rectangle $a \leq x(1) \leq b$, $c \leq x(2) \leq d$:

```

function cout = rectconstr(x,a,b,c,d)
% Negative when x is in the rectangle [a,b][c,d]
% First check that a,b,c,d are in the correct order

if (b <= a) || (d <= c)
    error('Give a rectangle a < b, c < d')
end

cout = max([(x(1)-b), (x(2)-d), (a-x(1)), (c-x(2))]);

```


Following the prescription of using the minimum of nonlinear constraint functions, for the L-shaped region, the nonlinear constraint function is:

```

function [c,ceq] = rectconstrfcn(x)

ceq = []; % no equality constraint
F(1) = rectconstr(x,-1,1,0,1); % one rectangle
F(2) = rectconstr(x,0,1,-1,1); % another rectangle
c = min(F); % for OR constraints

```



Code for creating the figure

Plot `rectconstrfcn` over the region $\max |x| \leq 2$ for $a = -1$, $b = 1$, $c = 0$, $d = 1$:

```
[xx,yy] = meshgrid(-2:.1:2);  
x = [xx(:),yy(:)]; % one row per point  
  
z = zeros(length(x),1); % allocate  
for ii = 1:length(x)  
    [z(ii),~] = rectconstrfcn(x(ii,:));  
end  
  
z = reshape(z,size(xx));
```

```
surf(xx,yy,z)
colorbar
axis equal
xlabel('x');ylabel('y')
view(0,90)
```

Suppose your objective function is

```
fun = @(x)exp(x(1)) * (4*x(1)^2 + 2*x(2)^2 +
4*x(1)*x(2) + 2*x(2) + 1);
```

Minimize fun over the L-shaped region:

```
opts = optimoptions(@fmincon,'Algorithm','interior-point','Display','off');
x0 = [-.5,.6]; % an arbitrary guess
[xsol,fval,eflag,output] = fmincon(fun,x0,[],[],[],[],[],[],@rectconstrfcn,
```

```
xsol =
```

```
    0.4998   -0.9996
```

```
fval =
```

```
    2.4649e-07
```

```
eflag =
```

```
    1
```

```
output =
```

```
    iterations: 17
    funcCount: 59
  constrviolation: 0
    stepsize: 1.8763e-04
    algorithm: 'interior-point'
```

```
firstorderopt: 4.9302e-07
cgiterations: 0
message: 'Local minimum found that satisfies the constraints.'
```

```
Optimization complet...'
```

Clearly, the solution `xsol` is inside the L-shaped region. The exit flag is 1, indicating that `xsol` is a local minimum.

How to Use All Types of Constraints

This section contains an example of a nonlinear minimization problem with all possible types of constraints. The objective function is in the local function `myobj(x)`. The nonlinear constraints are in the local function `myconstr(x)`. This example does not use gradients.

```
function [x fval exitflag] = fullexample
x0 = [1; 4; 5; 2; 5];
lb = [-Inf; -Inf; 0; -Inf; 1];
ub = [ Inf;  Inf; 20; Inf; Inf];
Aeq = [1 -0.3 0 0 0];
beq = 0;
A = [0 0 0 -1 0.1
     0 0 0 1 -0.5
     0 0 -1 0 0.9];
b = [0; 0; 0];
opts = optimoptions(@fmincon,'Algorithm','sqp');

[x,fval,exitflag]=fmincon(@myobj,x0,A,b,Aeq,beq,lb,ub,...
                        @myconstr,opts)

%-----
function f = myobj(x)

f = 6*x(2)*x(5) + 7*x(1)*x(3) + 3*x(2)^2;

%-----
function [c, ceq] = myconstr(x)

c = [x(1) - 0.2*x(2)*x(5) - 71
```

```
    0.9*x(3) - x(4)^2 - 67];  
ceq = 3*x(2)^2*x(5) + 3*x(1)^2*x(3) - 20.875;
```

Calling `fullexample` produces the following display in the Command Window:

```
[x fval exitflag] = fullexample;
```

```
Local minimum found that satisfies the constraints.
```

```
Optimization completed because the objective function is non-decreasing in  
feasible directions, to within the default value of the function tolerance,  
and constraints are satisfied to within the default value of the constraint tolerance.
```

```
x =
```

```
    0.6114  
    2.0380  
    1.3948  
    0.1572  
    1.5498
```

```
fval =
```

```
    37.3806
```

```
exitflag =
```

```
    1
```

Objective and Nonlinear Constraints in the Same Function

This example shows how to avoid calling a function twice when it computes values for both objective and constraints.

You typically use such a function in a simulation. Solvers such as `fmincon` evaluate the objective and nonlinear constraint functions separately. This evaluation is wasteful when you use the same calculation for both results.

To avoid wasting time, have your calculation use a nested function to evaluate the objective and constraint functions only when needed, by retaining the values of time-consuming calculations. Using a nested function avoids using global variables, yet lets intermediate results be retained and shared between the objective and constraint functions.

Step 1. Function that computes objective and constraints.

For example, suppose `computeall` is the expensive (time-consuming) function called by both the objective function and by the nonlinear constraint functions. Suppose you want to use `fmincon` as your optimizer.

Write a function that computes a portion of Rosenbrock's function `f1` and a nonlinear constraint `c1` that keeps the solution in a disk of radius 1 around the origin. Rosenbrock's function is

$$f(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2,$$

which has a unique minimum value of 0 at (1,1). See "Solve a Constrained Nonlinear Problem" on page 1-3.

In this example there is no nonlinear equality constraint, so `ceq1 = []`. Add a `pause(1)` statement to simulate an expensive computation.

```
function [f1,c1,ceq1] = computeall(x)
    ceq1 = [];
    c1 = norm(x)^2 - 1;
    f1 = 100*(x(2) - x(1)^2)^2 + (1-x(1))^2;
    pause(1) % simulate expensive computation
end
```


Save `computeall.m` as a file on your MATLAB path.

Step 2. Embed function in nested function that keeps recent values.

Suppose the objective function is

$$y = 100(x_2 - x_1^2)^2 + (1 - x_1)^2 + 20*(x_3 - x_4^2)^2 + 5*(1 - x_4)^2.$$

`computeall` returns the first part of the objective function. Embed the call to `computeall` in a nested function:

```
function [x,f,eflag,outpt] = runobjconstr(x0,opts)

if nargin == 1 % No options supplied
    opts = [];
end

xLast = []; % Last place computeall was called
myf = []; % Use for objective at xLast
myc = []; % Use for nonlinear inequality constraint
myceq = []; % Use for nonlinear equality constraint

fun = @objfun; % the objective function, nested below
cfun = @constr; % the constraint function, nested below

% Call fmincon
[x,f,eflag,outpt] = fmincon(fun,x0,[],[],[],[],[],[],cfun,opts);

function y = objfun(x)
    if ~isequal(x,xLast) % Check if computation is necessary
        [myf,myc,myceq] = computeall(x);
        xLast = x;
    end
    % Now compute objective function
    y = myf + 20*(x(3) - x(4)^2)^2 + 5*(1 - x(4))^2;
end

function [c,ceq] = constr(x)
    if ~isequal(x,xLast) % Check if computation is necessary
```

```
        [myf,myc,myceq] = computeall(x);
        xLast = x;
    end
    % Now compute constraint functions
    c = myc; % In this case, the computation is trivial
    ceq = myceq;
end

end
```

Save the nested function as a file named `runobjconstr.m` on your MATLAB path.

Step 3. Time to run with nested function.

Run the file, timing the call with `tic` and `toc`.

```
opts = optimoptions(@fmincon,'Algorithm','interior-point','Display','off');
x0 = [-1,1,1,2];
tic
[x,fval,exitflag,output] = runobjconstr(x0,opts);
toc
```

Elapsed time is 203.797275 seconds.

Step 4. Time to run without nested function.

Compare the times to run the solver with and without the nested function. For the run without the nested function, save `myrosen2.m` as the objective function file, and `constr.m` as the constraint:

```
function y = myrosen2(x)
    f1 = computeall(x); % get first part of objective
    y = f1 + 20*(x(3) - x(4)^2)^2 + 5*(1 - x(4))^2;
end

function [c,ceq] = constr(x)
    [~,c,ceq] = computeall(x);
end
```

Run `fmincon`, timing the call with `tic` and `toc`.

```
tic
[x,fval,exitflag,output] = fmincon(@myrosen2,x0,...
                                [],[],[],[],[],[],[],@constr,opts);
toc
```

Elapsed time is 406.771978 seconds.

The solver takes twice as long as before, because it evaluates the objective and constraint separately.

Step 5. Save computing time with parallel computing.

If you have a Parallel Computing Toolbox license, you can save even more time by setting the `UseParallel` option to `true`.

```
parpool
```

```
Starting parallel pool (parpool) using the 'local' profile ... connected to
```

```
ans =
```

```
Pool with properties:
```

```
    Connected: true
   NumWorkers: 4
    Cluster: local
 AttachedFiles: {}
 IdleTimeout: 30 minute(s) (30 minutes remaining)
  SpmdEnabled: true
```

```
opts = optimoptions(opts,'UseParallel',true);
tic
[x,fval,exitflag,output] = runobjconstr(x0,opts);
toc
```

Elapsed time is 97.528110 seconds.

In this case, enabling parallel computing cuts the computational time in half.

Compare the runs with parallel computing, with and without a nested function:

```
tic
[x,fval,exitflag,output] = fmincon(@myrosen2,x0,...
                                [],[],[],[],[],[],[],@constr,opts);
toc
```

Elapsed time is 188.985178 seconds.

In this example, computing in parallel but not nested takes about the same time as computing nested but not parallel. Computing both nested and parallel takes half the time of using either alone.

Related Examples

- “Solve a Constrained Nonlinear Problem” on page 1-3

Concepts

- “Optimizing a Simulation or Ordinary Differential Equation” on page 4-33
- “Parallel Computing”

Passing Extra Parameters

Extra Parameters, Fixed Variables, or Data

Sometimes objective or constraint functions have parameters in addition to the independent variable. The extra parameters can be data, or can represent variables that do not change during the optimization. There are three methods of passing these parameters:

- “Anonymous Functions” on page 2-53
- “Nested Functions” on page 2-55
- “Global Variables” on page 2-56

Global variables are troublesome because they do not allow names to be reused among functions. It is better to use one of the other two methods.

For example, suppose you want to minimize the function

$$f(x) = \left(a - bx_1^2 + x_1^{4/3}\right)x_1^2 + x_1x_2 + \left(-c + cx_2^2\right)x_2^2 \quad (2-1)$$

for different values of a , b , and c . Solvers accept objective functions that depend only on a single variable (x in this case). The following sections show how to provide the additional parameters a , b , and c . The solutions are for parameter values $a = 4$, $b = 2.1$, and $c = 4$ near $x_0 = [0.5 \ 0.5]$ using `fminunc`.

Anonymous Functions

To pass parameters using anonymous functions:

- 1 Write a 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;
```

- 2 Assign values to the parameters and define a function handle `f` to an anonymous function by entering the following commands at the MATLAB prompt:

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

3 Call the solver `fminunc` with the anonymous function:

```
[x,fval] = fminunc(f,x0)
```

The following output is displayed in the command window:

```
Local minimum found.
```

```
Optimization completed because the size of the gradient is less than
the default value of the function tolerance.
```

```
x =
   -0.0898    0.7127
```

```
fval =
   -1.0316
```

Note The parameters passed in the anonymous function are those that exist at the time the anonymous function is created. Consider the example

```
a = 4; b = 2.1; c = 4;
f = @(x)parameterfun(x,a,b,c)
```

Suppose you subsequently change, `a` to 3 and run

```
[x,fval] = fminunc(f,x0)
```

You get the same answer as before, since `parameterfun` uses `a = 4`, the value when `f` was created.

To change the parameters that are passed to the function, renew the anonymous function by reentering it:

```
a = 3;
f = @(x)parameterfun(x,a,b,c)
```

You can create anonymous functions of more than one argument. For example, to use `lsqcurvefit`, first create a function that takes two input arguments, `x` and `xdata`:

```
fh = @(x,xdata)(sin(x).*xdata +(x.^2).*cos(xdata));
x = pi; xdata = pi*[4;2;3];
fh(x, xdata)
```

```
ans =
```

```
    9.8696
    9.8696
   -9.8696
```

Now call `lsqcurvefit`:

```
% Assume ydata exists
x = lsqcurvefit(fh,x,xdata,ydata)
```

Nested Functions

To pass the parameters for Equation 2-1 via a nested function, write a single file that

- Accepts `a`, `b`, `c`, and `x0` as inputs
- Contains the objective function as a nested function
- Calls `fminunc`

Here is the code for the function file for this example:

```
function [x,fval] = runnested(a,b,c,x0)
[x,fval] = fminunc(@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
```

The objective function is the nested function `nestedfun`, which has access to the variables `a`, `b`, and `c`.

To run the optimization, enter:

```
a = 4; b = 2.1; c = 4;% Assign parameter values
x0 = [0.5,0.5];
[x,fval] = runnested(a,b,c,x0)
```

The output is the same as in “Anonymous Functions” on page 2-53.

Global Variables

Global variables can be troublesome, so it is better to avoid using them. To use global variables, declare the variables to be global in the workspace and in the functions that use the variables.

1 Write a function file:

```
function y = globalfun(x)
global 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;
```

2 In your MATLAB workspace, define the variables and run `fminunc`:

```
global a b c;
a = 4; b = 2.1; c = 4; % Assign parameter values
x0 = [0.5,0.5];
[x,fval] = fminunc(@globalfun,x0)
```

The output is the same as in “Anonymous Functions” on page 2-53.

What Are Options?

Options are a way of combining a set of name-value pairs. They are useful because they allow you to:

- Tune or modify the optimization process.
- Select extra features, such as output functions and plot functions.
- Save and reuse settings.

They simplify solver syntax—you don't have to include a lot of name-value pairs in a call to a solver.

To see how to set and change options, see “Set and Change Options” on page 2-59.

For an overview of all options, including which solvers use each option, see “Optimization Options Reference” on page 9-7.

Options in Common Use: Tuning and Troubleshooting

You set or change options when the default settings do not work sufficiently well. This can mean the solver takes too long to converge, the solver fails, or you are unsure of the reliability of the result.

To tune your solver for improved speed or accuracy, try setting these options first:

- “Choosing the Algorithm” on page 2-7 — `Algorithm`
- “Tolerances and Stopping Criteria” on page 2-65 — `TolFun`, `TolX`, `MaxFunEvals`, and `MaxIter`
- Finite differences — `FinDiffType` and `FinDiffRelStep`

To diagnose and troubleshoot, try setting these options first:

- “Iterative Display” on page 3-17 — `Display`
- Function evaluation errors — `FunValCheck`
- “Plot Functions” on page 3-31 and “Output Functions” on page 3-37 — `PlotFcns` and `OutputFunction`

Related Examples

- “Improve Results”

Concepts

- “Solver Outputs and Iterative Display”

Set and Change Options

The recommended way to set options is to use the `optimoptions` function. For example, to set the `fmincon` algorithm to `sqp`, set iterative display, and set a small value of the `TolCon` tolerance:

```
options = optimoptions('fmincon',...  
    'Algorithm','sqp','Display','iter','TolCon',1e-12);
```

Note Use `optimset` instead of `optimoptions` for the `fminbnd`, `fminsearch`, `fzero`, and `lsqnonneg` solvers. These are the solvers that do not require an Optimization Toolbox license.

Change options as follows:

- Dot notation. For example,

```
options.TolX = 1e-10;
```

- `optimoptions`. For example,

```
options = optimoptions(options,'TolX',1e-10);
```

Ensure that you pass `options` in your solver call. For example,

```
[x,fval] =  
fmincon(@objfun,x0,[],[],[],[],lb,ub,@nonlcon,options);
```

You can also set and change options using the “Optimization App” on page 5-2.

Choose Between `optimoptions` and `optimset`

Previously, the recommended way to set options was to use `optimset`. Now the general recommendation is to use `optimoptions`, with some caveats listed below.

`optimset` still works, and it is the only way to set options for solvers that are available without an Optimization Toolbox license: `fminbnd`, `fminsearch`, `fzero`, and `lsqnonneg`.

Note Some other toolboxes use optimization options and require you to pass in options created using `optimset`, not `optimoptions`. Check the documentation for your toolboxes.

`optimoptions` organizes options by solver, with a more focused and comprehensive display than `optimset`:

- Creates and modifies only the options that apply to a solver
- Shows your option choices and default values for a specific solver/algorithm
- Displays links for more information on solver options and other available solver algorithms

`intlinprog` uses only `optimoptions` options.

The main difference in creating options is:

- For `optimoptions`, you include the solver name as the first argument.

```
options = optimoptions(SolverName,Name,Value,...)
```

- For `optimset`, the syntax does not include the solver name.

```
options = optimset(Name,Value,...)
```

In both cases, you can query or change options by using dot notation. See “Set and Change Options” on page 2-59 and “View Options” on page 2-63.

For example, compare the display of `optimoptions` to that of `optimset`.

```
options = optimset('GradObj','on')  
options = optimoptions(@fminunc,'GradObj','on')
```

options =

[fminunc](#) options:

Options used by current Algorithm ('trust-region'):
([Other available algorithms](#): 'quasi-newton')

Set by user:

GradObj: 'on'

Default:

Algorithm: 'trust-region'
DerivativeCheck: 'off'
Diagnostics: 'off'
DiffMaxChange: Inf
DiffMinChange: 0
Display: 'final'
FinDiffRelStep: 1.4901e-08
FinDiffType: 'forward'
FunValCheck: 'off'
HessMult: []
HessPattern: 'sparse(ones(numberOfVariables))'
Hessian: 'off'
MaxFunEvals: '100*numberOfVariables'
MaxIter: 400
MaxPCGIter: 'max(1,floor(numberOfVariables/2))'
OutputFcn: []
PlotFcns: []
PrecondBandWidth: 0
TolFun: 1.0000e-06
TolPCG: 0.1000
TolX: 1.0000e-06
TypicalX: 'ones(numberOfVariables,1)'

[Show options](#) not used by current Algorithm ('trust-region')

options =

View Options

You can view the value of a particular option by using dot notation. For example,

```
options = optimoptions('fmincon','Algorithm','interior-point');
```

To view the value of the TolX tolerance:

```
options.TolX
```

```
ans =
```

```
1.0000e-10
```

You can view the value of all options, and see which ones have nondefault values, by entering the options name:

```
options
```

```
options =
```

```
fmincon options:
```

```
Options used by current Algorithm ('interior-point'):
```

```
(Other available algorithms: 'active-set', 'sqp', 'trust-region-reflective')
```

```
Set by user:
```

```
Algorithm: 'interior-point'
```

```
Default:
```

```
AlwaysHonorConstraints: 'bounds'
```

```
DerivativeCheck: 'off'
```

```
Diagnostics: 'off'
```

```
DiffMaxChange: Inf
```

```
DiffMinChange: 0
```

```
Display: 'final'
```

```
FinDiffRelStep: 1.4901e-08
```

```
FinDiffType: 'forward'
```

```
FunValCheck: 'off'
```

```
GradConstr: 'off'
```

```
GradObj: 'off'
```

```
HessFcn: []
```

```
HessMult: []
```

```
Hessian: 'bfgs'
```

```
InitBarrierParam: 0.1000
```

```
InitTrustRegionRadius: 'sqrt(numberOfVariables)'
```

```
MaxFunEvals: 3000
```

```
MaxIter: 1000
```

```
MaxProjCGIter: '2*(numberOfVariables-numberOfEqualities)'
```

```
ObjectiveLimit: -1.0000e+20
```

```
OutputFcn: []
```

```
PlotFcns: []
```

```
ScaleProblem: 'none'
```

```
SubproblemAlgorithm: 'ldl-factorization'
```

```
TolCon: 1.0000e-06
```

```
TolFun: 1.0000e-06
```

```
TolProjCG: 0.0100
```

```
TolProjCGAbs: 1.0000e-10
```

```
TolX: 1.0000e-10
```

```
TypicalX: 'ones(numberOfVariables,1)'
```

```
UseParallel: 'never'
```

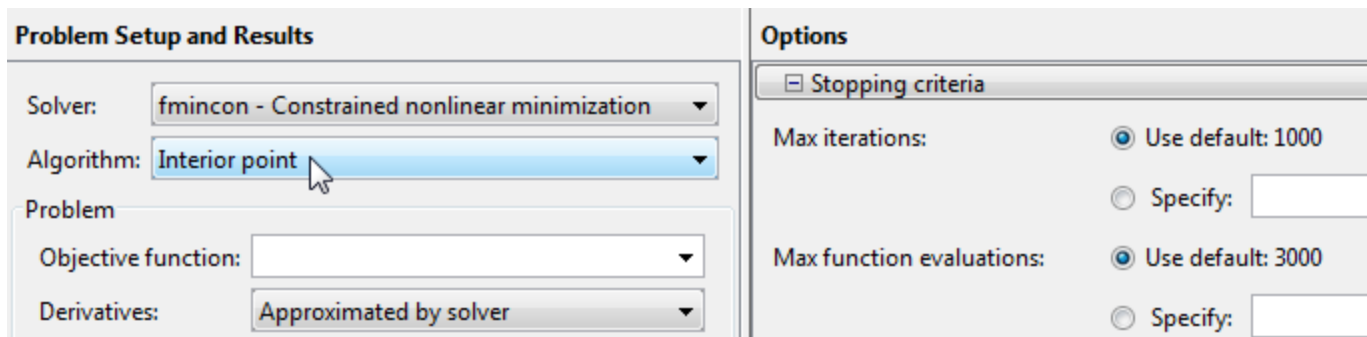

Tolerances and Stopping Criteria

The number of iterations in an optimization depends on a solver's *stopping criteria*. These criteria include several *tolerances* you can set. Generally, a tolerance is a threshold which, if crossed, stops the iterations of a solver.

Set tolerances and other criteria using `optimoptions` as explained in “Set and Change Options” on page 2-59.

Tip Generally set tolerances such as `TolFun` and `TolX` to be well above `eps`, and usually above $1e-14$. Setting small tolerances does not always result in accurate results. Instead, a solver can fail to recognize when it has converged, and can continue futile iterations. A tolerance value smaller than `eps` effectively disables that stopping condition.

You can find the default tolerances in the “Optimization App” on page 5-2. Some default tolerances differ for different algorithms, so set both the solver and the algorithm.



`optimoptions` displays default tolerances. For example,

```
options = optimoptions('fmincon')
```

options =

fmincon options:

Options used by current Algorithm ('interior-point'):

(Other available algorithms: 'active-set', 'sqp', 'trust-region-reflective')

Set by user:

No options set by user.

Default:

```
          Algorithm: 'interior-point'
AlwaysHonorConstraints: 'bounds'
  DerivativeCheck: 'off'
    Diagnostics: 'off'
  DiffMaxChange: Inf
  DiffMinChange: 0
    Display: 'final'
  FinDiffRelStep: 'sqrt(eps)'
  FinDiffType: 'forward'
  FunValCheck: 'off'
    GradConstr: 'off'
      GradObj: 'off'
      HessFcn: []
      Hessian: 'bfgs'
      HessMult: []
  InitBarrierParam: 0.1000
  InitTrustRegionRadius: 'sqrt(numberOfVariables)'
```

MaxFunEvals: 3000

MaxIter: 1000

MaxProjCGIter: '2*(numberOfVariables-numberOfEqualities)'

ObjectiveLimit: -1.0000e+20

OutputFcn: []

PlotFcns: []

ScaleProblem: 'none'

SubproblemAlgorithm: 'ldl-factorization'

TolCon: 1.0000e-06

TolFun: 1.0000e-06


TolProjCG: 0.0100

TolProjCGAbs: 1.0000e-10

TolX: 1.0000e-10

TypicalX: ones(numberOfVariables, 1)'

- TolX is a lower bound on the size of a step, meaning the norm of $(x_i - x_{i+1})$. If the solver attempts to take a step that is smaller than TolX, the iterations end. TolX is sometimes used as a *relative* bound, meaning iterations end when $|x_i - x_{i+1}| < \text{TolX} * (1 + |x_i|)$, or a similar relative measure.



- TolFun is a lower bound on the change in the value of the objective function during a step. If $|f(x_i) - f(x_{i+1})| < \text{TolFun}$, the iterations end. TolFun is sometimes used as a *relative* bound, meaning iterations end when $|f(x_i) - f(x_{i+1})| < \text{TolFun} * (1 + |f(x_i)|)$, or a similar relative measure.
- TolFun is also a bound on the first-order optimality measure. If the optimality measure is less than TolFun, the iterations end. TolFun can also be a relative bound. First-order optimality measure is defined in “First-Order Optimality Measure” on page 3-12.
- TolCon is an upper bound on the magnitude of any constraint functions. If a solver returns a point x with $c(x) > \text{TolCon}$ or $|ceq(x)| > \text{TolCon}$, the solver reports that the constraints are violated at x . TolCon can also be a relative bound.

Note TolCon operates differently from other tolerances. If TolCon is not satisfied (i.e., if the magnitude of the constraint function exceeds TolCon), the solver attempts to continue, unless it is halted for another reason. A solver does not halt simply because TolCon is satisfied.

- `MaxIter` is a bound on the number of solver iterations. `MaxFunEvals` is a bound on the number of function evaluations. Iterations and function evaluations are discussed in “Iterations and Function Counts” on page 3-11.

There are two other tolerances that apply to particular solvers: `TolPCG` and `MaxPCGIter`. These relate to preconditioned conjugate gradient steps. For more information, see “Preconditioned Conjugate Gradient Method” on page 6-29.

There are several tolerances that apply only to the `fmincon` interior-point algorithm. For more information, see “Interior-Point Algorithm” on page 10-60.

There are several tolerances that apply only to `intlinprog`. See “Some “Integer” Solutions Are Not Integers” on page 6-167 and “Branch and Bound” on page 6-161.

Checking Validity of Gradients or Jacobians

In this section...

“How to Check Derivatives” on page 2-69

“Example: Checking Derivatives of Objective and Constraint Functions” on page 2-70

Many solvers allow you to supply a function that calculates first derivatives (gradients or Jacobians) of objective or constraint functions. You can check whether the derivatives calculated by your function match finite-difference approximations. This check can help you diagnose whether your derivative function is correct.

- If a component of the gradient function is less than 1, “match” means the absolute difference of the gradient function and the finite-difference approximation of that component is less than $1e-6$.
- Otherwise, “match” means that the relative difference is less than $1e-6$.

The `DerivativeCheck` option causes the solver to check the supplied derivative against a finite-difference approximation at just one point. If the finite-difference and supplied derivatives do not match, the solver errors. If the derivatives match to within $1e-6$, the solver reports the calculated differences, and continues iterating without further derivative checks. Solvers check the match at a point that is a small random perturbation of the initial point x_0 , modified to be within any bounds. Solvers do not include the computations for `DerivativeCheck` in the function count; see “Iterations and Function Counts” on page 3-11.

How to Check Derivatives

- At the MATLAB command line:
 - 1 Set the `GradObj`, `GradConstr`, or `Jacobian` options to 'on' with `optimoptions`. Make sure your objective or constraint functions supply the appropriate derivatives.
 - 2 Set the `DerivativeCheck` option to 'on'.

- Using the Optimization app:
 - 1 In the **Problem Setup and Results** pane, choose **Derivatives: Objective function: Gradient supplied** or **Nonlinear constraint function: Derivatives: Gradient supplied**. Make sure your objective or constraint functions supply the appropriate derivatives.
 - 2 In the **Options** pane, check **User-supplied derivatives > Validate user-supplied derivatives**

Central finite differences are more accurate than the default forward finite differences. To use central finite differences:

- At the MATLAB command line, set `FinDiffType` option to 'central' with `optimoptions`.
- Using the Optimization app, in the **Approximated derivatives** pane, set **Type** to central differences.

Example: Checking Derivatives of Objective and Constraint Functions

- “Objective and Constraint Functions” on page 2-70
- “Checking Derivatives at the Command Line” on page 2-71
- “Checking Derivatives with the Optimization App” on page 2-73

Objective and Constraint Functions

Consider the problem of minimizing the Rosenbrock function within the unit disk as described in “Solve a Constrained Nonlinear Problem” on page 1-3. The `rosenboth` function calculates the objective function and its gradient:

```
function [f g H] = rosenboth(x)

f = 100*(x(2) - x(1)^2)^2 + (1-x(1))^2;

if nargin > 1
    g = [-400*(x(2)-x(1)^2)*x(1)-2*(1-x(1));
         200*(x(2)-x(1)^2)];
```

```

        if nargout > 2
            H = [1200*x(1)^2-400*x(2)+2, -400*x(1);
                -400*x(1), 200];
        end
    end
end

```

rosenboth calculates the Hessian, too, but this example does not use the Hessian.

The `unitdisk2` function correctly calculates the constraint function and its gradient:

```

function [c,ceq,gc,gceq] = unitdisk2(x)
c = x(1)^2 + x(2)^2 - 1;
ceq = [ ];

if nargout > 2
    gc = [2*x(1);2*x(2)];
    gceq = [];
end

```

The `unitdiskb` function incorrectly calculates gradient of the constraint function:

```

function [c ceq gc gceq] = unitdiskb(x)
c = x(1)^2 + x(2)^2 - 1;
ceq = [ ];

if nargout > 2
    gc = [x(1);x(2)]; % Gradient incorrect: off by a factor of 2
    gceq = [];
end

```

Checking Derivatives at the Command Line

- 1 Set the options to use the interior-point algorithm, gradient of objective and constraint functions, and the `DerivativeCheck` option:

```

% For reproducibility--DerivativeCheck randomly perturbs the initial point
rng(0,'twister');
options = optimoptions(@fmincon,'Algorithm','interior-point',...

```

```
'DerivativeCheck','on','GradObj','on','GradConstr','on');
```

- 2 Solve the minimization with `fmincon` using the erroneous `unitdiskb` constraint function:

```
[x fval exitflag output] = fmincon(@rosenboth,...  
    [-1;2],[],[],[],[],[],[],[],@unitdiskb,options);
```

Derivative Check Information

Objective function derivatives:

Maximum relative difference between user-supplied
and finite-difference derivatives = 1.84768e-008.

Nonlinear inequality constraint derivatives:

Maximum relative difference between user-supplied
and finite-difference derivatives = 1.

User-supplied constraint derivative element (2,1): 1.99838

Finite-difference constraint derivative element (2,1): 3.99675

Error using `validateFirstDerivatives`

Derivative Check failed:

User-supplied and forward finite-difference derivatives
do not match within 1e-006 relative tolerance.

Error in `fmincon` at 805

```
validateFirstDerivatives(funfcn,confcn,X, ...
```

The constraint function does not match the calculated gradient, encouraging
you to check the function for an error.

- 3 Replace the `unitdiskb` constraint function with `unitdisk2` and run the minimization again:

```
[x fval exitflag output] = fmincon(@rosenboth,...  
    [-1;2],[],[],[],[],[],[],[],@unitdisk2,options);
```

Derivative Check Information

Objective function derivatives:
Maximum relative difference between user-supplied
and finite-difference derivatives = 1.28553e-008.

Nonlinear inequality constraint derivatives:
Maximum relative difference between user-supplied
and finite-difference derivatives = 1.46443e-008.

Derivative Check successfully passed.

Local minimum found that satisfies the constraints...

Checking Derivatives with the Optimization App

To set up the example using correct derivative functions, but starting from [0 0], using the Optimization app:

- 1** Launch the Optimization app by entering `optimtool` at the command line.
- 2** Set the **Problem Setup and Results** pane to match the following figure:

Problem Setup and Results

Solver:

Algorithm:

Problem

Objective function:

Derivatives:

Start point:

Constraints:

Linear inequalities: A: b:

Linear equalities: Aeq: beq:

Bounds: Lower: Upper:

Nonlinear constraint function:

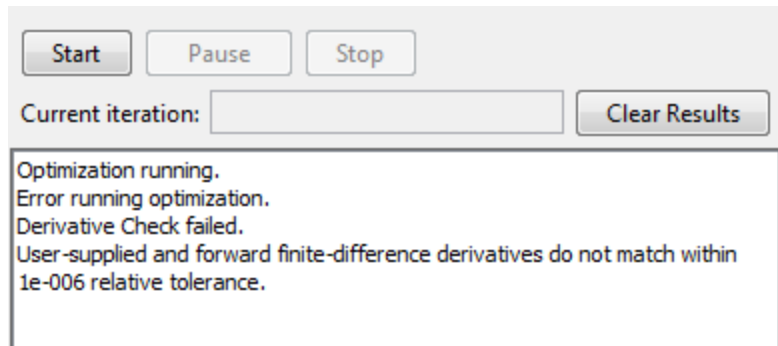
Derivatives:

3 Set the **Options** pane to match the following figure:

The image shows a software interface with two main sections: 'User-supplied derivatives' and 'Approximated derivatives'.
The 'User-supplied derivatives' section has a checked checkbox for 'Validate user-supplied derivatives'. Below it, 'Hessian sparsity pattern' has a selected radio button for 'Use default: sparse(ones(numberOfVariables))' and an unselected one for 'Specify:' with an empty text box. 'Hessian multiply function' has a selected radio button for 'Use default: No multiply function' and an unselected one for 'Specify:' with an empty text box.
The 'Approximated derivatives' section has a dropdown menu for 'Finite differences $f(x + r*x) - f(x)$ ' set to 'forward differences'. 'Relative perturbation vector r:' has a selected radio button for 'Use default: sqrt(eps)*ones(numberOfVariables,1)' and an unselected one for 'Specify:' with an empty text box. 'Minimum perturbation |r*x|:' has a selected radio button for 'Use default: 0' and an unselected one for 'Specify:' with an empty text box. 'Maximum perturbation |r*x|:' has a selected radio button for 'Use default: Inf' and an unselected one for 'Specify:' with an empty text box. At the bottom, there is an unchecked checkbox for 'Evaluate in parallel'.


- 4 Press the **Start** button under **Run solver and view results**.

The output screen displays

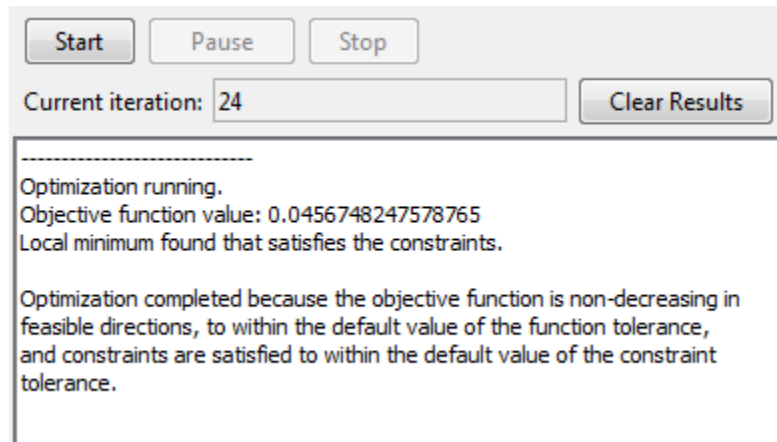


The forward finite difference approximation is inaccurate enough near $[0 \ 0]$ that the derivative check fails.

- 5 To use the more accurate central differences, select `central differences` in the **Approximated derivatives > Type** pane:



- 6 Click **Run solver and view results > Clear Results**, then **Start**. This time the derivative check is successful:



The derivative check also succeeds when you select the initial point $[-1 \ 2]$, or most random points.

Examining Results

- “Current Point and Function Value” on page 3-2
- “Exit Flags and Exit Messages” on page 3-3
- “Iterations and Function Counts” on page 3-11
- “First-Order Optimality Measure” on page 3-12
- “Iterative Display” on page 3-17
- “Output Structures” on page 3-26
- “Lagrange Multiplier Structures” on page 3-27
- “Hessian” on page 3-28
- “Plot Functions” on page 3-31
- “Output Functions” on page 3-37

Current Point and Function Value

The *current point* and *function value* are the first two outputs of all Optimization Toolbox solvers.

- The current point is the final point in the solver iterations. It is the best point the solver found in its run.
 - If you call a solver without assigning a value to the output, the default output, `ans`, is the current point.
- The function value is the value of the objective function at the current point.
 - The function value for least-squares solvers is the sum of squares, also known as the residual norm.
 - `fgoalattain`, `fminimax`, and `fsolve` return a vector function value.
 - Sometimes `fval` or `Fval` denotes function value.

Exit Flags and Exit Messages

In this section...
“Exit Flags” on page 3-3
“Exit Messages” on page 3-5
“Enhanced Exit Messages” on page 3-5
“Exit Message Options” on page 3-9

Exit Flags

When an optimization solver completes its task, it sets an *exit flag*. An exit flag is an integer that is a code for the reason the solver halted its iterations. In general:

- Positive exit flags correspond to successful outcomes.
- Negative exit flags correspond to unsuccessful outcomes.
- The zero exit flag corresponds to the solver being halted by exceeding an iteration limit or limit on the number of function evaluations (see “Iterations and Function Counts” on page 3-11, and also see “Tolerances and Stopping Criteria” on page 2-65).

A table of solver outputs in the solver’s function reference section lists the meaning of each solver’s exit flags.

Note Exit flags are not infallible guides to the quality of a solution. Many other factors, such as tolerance settings, can affect whether a solution is satisfactory to you. You are responsible for deciding whether a solver returns a satisfactory answer. Sometimes a negative exit flag does not correspond to a “bad” solution. Similarly, sometimes a positive exit flag does not correspond to a “good” solution.

You obtain an exit flag by calling a solver with the `exitflag` syntax. This syntax depends on the solver. For details, see the solver function reference pages. For example, for `fsolve`, the calling syntax to obtain an exit flag is

```
[x,fval,exitflag] = fsolve(...)
```

The following example uses this syntax. Suppose you want to solve the system of nonlinear equations

$$\begin{aligned}2x_1 - x_2 &= e^{-x_1} \\ -x_1 + 2x_2 &= e^{-x_2}.\end{aligned}$$

Write these equations as an anonymous function that gives a zero vector at a solution:

```
myfcn = @(x)[2*x(1) - x(2) - exp(-x(1));
            -x(1) + 2*x(2) - exp(-x(2))];
```

Call `fsolve` with the `exitflag` syntax at the initial point `[-5 -5]`:

```
[xfinal fval exitflag] = fsolve(myfcn,[-5 -5])
```

Equation solved.

`fsolve` completed because the vector of function values is near zero as measured by the default value of the function tolerance, and the problem appears regular as measured by the gradient.

```
xfinal =
    0.5671    0.5671
```

```
fval =
    1.0e-006 *
   -0.4059
   -0.4059
```

```
exitflag =
    1
```

In the table for `fsolve` under “Output Arguments” on page 10-97, you find that an exit flag value 1 means “Function converged to a solution x .” In other words, `fsolve` reports `myfcn` is nearly zero at $x = [0.5671 \ 0.5671]$.

Exit Messages

Each solver issues a message to the MATLAB command window at the end of its iterations. This message explains briefly why the solver halted. The message might give more detail than the exit flag.

Many examples in this documentation show exit messages. For example, see “Minimizing at the Command Line” on page 1-10, or “Step 2: Invoke one of the unconstrained optimization routines.” on page 6-17. The example in the previous section, “Exit Flags” on page 3-3, shows the following exit message:

```
Equation solved.
```

```
fsolve completed because the vector of function values is near
zero as measured by the default value of the function tolerance,
and the problem appears regular as measured by the gradient.
```

This message is more informative than the exit flag. The message indicates that the gradient is relevant. The message also states that the function tolerance controls how near 0 the vector of function values must be for `fsolve` to regard the solution as completed.

Enhanced Exit Messages

Some solvers have exit messages that contain links for more information. There are two types of links:

- Links on words or phrases. If you click such a link, a window opens that displays a definition of the term, or gives other information. The new window can contain links to the Help browser documentation for more detailed information.
- A link as the last line of the display saying `<stopping criteria details>`. If you click this link, MATLAB displays more detail about the reason the solver halted.

The `fminunc` solver has enhanced exit messages:

```
opts = optimoptions(@fminunc,'Algorithm','quasi-newton');
% 'trust-region' needs gradient
[xfinal fval exitflag] = fminunc(@sin,0,opts)
```

This yields the following results:

[Local minimum found.](#)

Optimization completed because the [size of the gradient](#) is less than the default value of the [function tolerance](#).

[<stopping criteria details>](#)

```
xfinal =  
  
    -1.5708  
  
fval =  
  
    -1.0000  
  
exitflag =  
  
     1
```


Each of the underlined words or phrases contains a link that provides more information.

- The [<stopping criteria details>](#) link prints the following to the MATLAB command line:

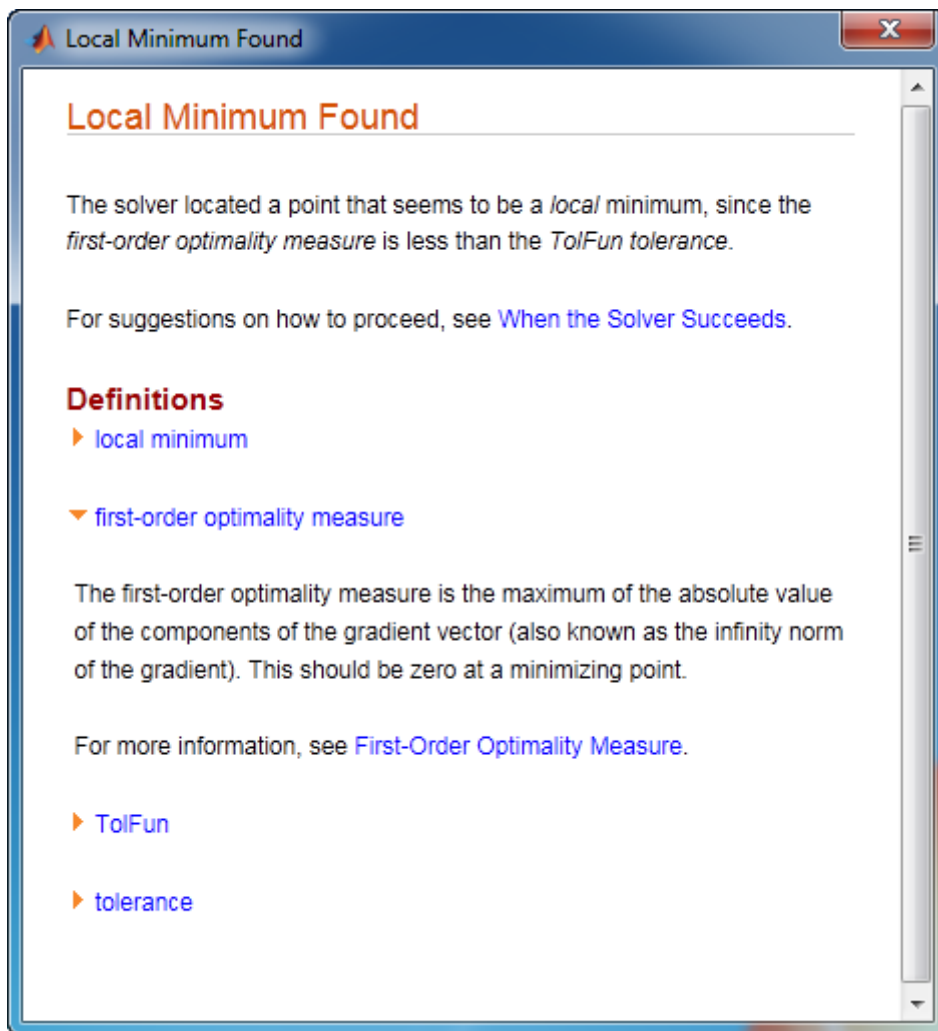
```
Optimization completed: The first-order optimality measure, 0.000000e+000, is less  
than the default value of options.TolFun = 1.000000e-006.
```

Optimization Metric	User Options
relative norm(gradient) = 0.00e+000	TolFun = 1e-006 (default)

- The other links bring up a help window with term definitions. For example, clicking the [Local minimum found](#) link opens the following window:



Clicking the first-order optimality measure expander link brings up the definition of first-order optimality measure for `fminunc`:



The expander link is a way to obtain more information in the same window. Clicking the `first-order optimality measure` expander link again closes the definition.

- The other links open the Help Viewer.

Exit Message Options

Set the `Display` option to control the appearance of both exit messages and iterative display. For more information, see “Iterative Display” on page 3-17. The following table shows the effect of the various settings of the `Display` option.

Value of the Display Option	Output to Command Window	
	Exit message	Iterative Display
'none', or the synonymous 'off'	None	None
'final' (default for most solvers)	Default	None
'final-detailed'	Detailed	None
'iter'	Default	Yes
'iter-detailed'	Detailed	Yes
'notify'	Default only if <code>exitflag ≤ 0</code>	None
'notify-detailed'	Detailed only if <code>exitflag ≤ 0</code>	None

For example,

```
opts = optimoptions(@fminunc,'Display','iter-detailed','Algorithm','quasi-n
[xfinal fval] = fminunc(@cos,1,opts);
```

yields the following display:

```
>> opts = optimoptions(@fminunc,'Display','iter-detailed','Algorithm','quasi-newton')
[xfinal fval] = fminunc(@cos,1,opts);
```

Iteration	Func-count	f(x)	Step-size	First-order optimality
0	2	0.540302		0.841
1	6	-0.990628	2.38223	0.137
2	10	-1	0.351894	0.000328
3	12	-1	1	1.03e-06

Optimization completed: The [first-order optimality measure](#), 5.602276e-07, is less than `options.TolFun` = 1.000000e-06.

Optimization Metric	Options
relative norm(gradient) = 5.60e-07	TolFun = 1e-06 (default)

Iterations and Function Counts

In general, Optimization Toolbox solvers iterate to find an optimum. This means a solver begins at an initial value x_0 , performs some intermediate calculations that eventually lead to a new point x_1 , and then repeats the process to find successive approximations x_2, x_3, \dots of the local minimum. Processing stops after some number of iterations k .

At any step, intermediate calculations may involve evaluating the objective function and constraints, if any, at points near the current iterate x_i . For example, the solver may estimate a gradient by finite differences. At each of these nearby points, the function count (F-count) is increased by one.

- If there are no constraints, the F-count reports the total number of objective function evaluations.
- If there are constraints, the F-count reports only the number of points where function evaluations took place, not the total number of evaluations of constraint functions.
- If there are many constraints, the F-count can be significantly less than the total number of function evaluations.

F-count is a header in the iterative display for many solvers. For an example, see “Interpreting the Result” on page 1-11.

The F-count appears in the output structure as `output.funcCount`. This enables you to access the evaluation count programmatically. For more information on output structures, see “Output Structures” on page 3-26.

Sometimes a solver attempts a step, and rejects the attempt. The `trust-region`, `trust-region-reflective`, and `trust-region-dogleg` algorithms count these failed attempts as iterations, and report the (unchanged) result in the iterative display. The `interior-point`, `active-set`, and `levenberg-marquardt` algorithms do not count such an attempt as an iteration, and do not report the attempt in the iterative display. All attempted steps increase the F-count, regardless of the algorithm.

First-Order Optimality Measure

In this section...

“What Is First-Order Optimality Measure?” on page 3-12

“Stopping Rules Related to First-Order Optimality” on page 3-12

“Unconstrained Optimality” on page 3-13

“Constrained Optimality Theory” on page 3-13

“Constrained Optimality in Solver Form” on page 3-15

What Is First-Order Optimality Measure?

First-order optimality is a measure of how close a point x is to optimal. Most Optimization Toolbox solvers use this measure, though it has different definitions for different algorithms. First-order optimality is a necessary condition, but it is not a sufficient condition. In other words:

- The first-order optimality measure must be zero at a minimum.
- A point with first-order optimality equal to zero is not necessarily a minimum.

For general information about first-order optimality, see Nocedal and Wright [31]. For specifics about the first-order optimality measures for Optimization Toolbox solvers, see “Unconstrained Optimality” on page 3-13, “Constrained Optimality Theory” on page 3-13, and “Constrained Optimality in Solver Form” on page 3-15.

Stopping Rules Related to First-Order Optimality

The TolFun tolerance relates to the first-order optimality measure. Typically, if the first-order optimality measure is less than TolFun, solver iterations end.

Some solvers or algorithms use *relative* first-order optimality as a stopping criterion. Solver iterations end if the first-order optimality measure is less than μ times TolFun, where μ is either:

- The infinity norm (maximum) of the gradient of the objective function at x_0

- The infinity norm (maximum) of inputs to the solver, such as f or b in `linprog` or H in `quadprog`

A relative measure attempts to account for the scale of a problem. Multiplying an objective function by a very large or small number does not change the stopping condition for a relative stopping criterion, but does change it for an unscaled one.

Solvers with enhanced exit messages state, in the stopping criteria details, when they use relative first-order optimality.

Unconstrained Optimality

For a smooth unconstrained problem,

$$\min_x f(x),$$

the first-order optimality measure is the infinity norm (meaning maximum absolute value) of $\nabla f(x)$, which is:

$$\text{first-order optimality measure} = \max_i |(\nabla f(x))_i| = \|\nabla f(x)\|_\infty.$$

This measure of optimality is based on the familiar condition for a smooth function to achieve a minimum: its gradient must be zero. For unconstrained problems, when the first-order optimality measure is nearly zero, the objective function has gradient nearly zero, so the objective function could be near a minimum. If the first-order optimality measure is not small, the objective function is not minimal.

Constrained Optimality Theory

This section summarizes the theory behind the definition of first-order optimality measure for constrained problems. The definition as used in Optimization Toolbox functions is in “Constrained Optimality in Solver Form” on page 3-15.

For a smooth constrained problem, let g and h be vector functions representing all inequality and equality constraints respectively (meaning bound, linear, and nonlinear constraints):

$$\min_x f(x) \text{ subject to } g(x) \leq 0, h(x) = 0.$$

The meaning of first-order optimality in this case is more complex than for unconstrained problems. The definition is based on the Karush-Kuhn-Tucker (KKT) conditions. The KKT conditions are analogous to the condition that the gradient must be zero at a minimum, modified to take constraints into account. The difference is that the KKT conditions hold for constrained problems.

The KKT conditions use the auxiliary Lagrangian function:

$$L(x, \lambda) = f(x) + \sum \lambda_{g,i} g_i(x) + \sum \lambda_{h,i} h_i(x). \tag{3-1}$$

The vector λ , which is the concatenation of λ_g and λ_h , is the Lagrange multiplier vector. Its length is the total number of constraints.

The KKT conditions are:

$$\nabla_x L(x, \lambda) = 0, \tag{3-2}$$

$$\lambda_{g,i} g_i(x) = 0 \quad \forall i, \tag{3-3}$$

$$\begin{cases} g(x) \leq 0, \\ h(x) = 0, \\ \lambda_{g,i} \geq 0. \end{cases} \tag{3-4}$$

Solvers do not use the three expressions in Equation 3-4 in the calculation of optimality measure.

The optimality measure associated with Equation 3-2 is

$$\|\nabla_x L(x, \lambda)\| = \|\nabla f(x) + \sum \lambda_{g,i} \nabla g_i(x) + \sum \lambda_{h,i} \nabla h_{h,i}(x)\|. \tag{3-5}$$

The optimality measure associated with Equation 3-3 is

$$\|\overline{\lambda_g g(x)}\|, \quad (3-6)$$

where the norm in Equation 3-6 means infinity norm (maximum) of the vector $\overline{\lambda_{g,i} g_i(x)}$.

The combined optimality measure is the maximum of the values calculated in Equation 3-5 and Equation 3-6. Solvers that accept nonlinear constraint functions report constraint violations $g(x) > 0$ or $|h(x)| > 0$ as TolCon tolerance violations. See “Tolerances and Stopping Criteria” on page 2-65.

Constrained Optimality in Solver Form

Most constrained toolbox solvers separate their calculation of first-order optimality measure into bounds, linear functions, and nonlinear functions. The measure is the maximum of the following two norms, which correspond to Equation 3-5 and Equation 3-6:

$$\begin{aligned} \|\nabla_x L(x, \lambda)\| = & \|\nabla f(x) + A^T \lambda_{ineqlin} + Aeq^T \lambda_{eqlin} \\ & + \sum \lambda_{ineqnonlin,i} \nabla c_i(x) + \sum \lambda_{eqlin,i} \nabla ceq_i(x)\|, \end{aligned} \quad (3-7)$$

$$\|\overline{|l_i - x_i| \lambda_{lower,i}}, \overline{|x_i - u_i| \lambda_{upper,i}}, \overline{|Ax - b)_i| \lambda_{ineqlin,i}}, \overline{|c_i(x)| \lambda_{ineqnonlin,i}}\|, \quad (3-8)$$

where the norm of the vectors in Equation 3-7 and Equation 3-8 is the infinity norm (maximum). The subscripts on the Lagrange multipliers correspond to solver Lagrange multiplier structures. See “Lagrange Multiplier Structures” on page 3-27. The summations in Equation 3-7 range over all constraints. If a bound is $\pm\text{Inf}$, that term is not constrained, so it is not part of the summation.

Linear Equalities Only

For some large-scale problems with only linear equalities, the first-order optimality measure is the infinity norm of the *projected* gradient. In other words, the first-order optimality measure is the size of the gradient projected onto the null space of Aeq.

Bounded Least-Squares and Trust-Region-Reflective Solvers

For least-squares solvers and trust-region-reflective algorithms, in problems with bounds alone, the first-order optimality measure is the maximum over i of $|v_i * g_i|$. Here g_i is the i th component of the gradient, x is the current point, and

$$v_i = \begin{cases} |x_i - b_i| & \text{if the negative gradient points toward bound } b_i \\ 1 & \text{otherwise.} \end{cases}$$

If x_i is at a bound, v_i is zero. If x_i is not at a bound, then at a minimizing point the gradient g_i should be zero. Therefore the first-order optimality measure should be zero at a minimizing point.

Iterative Display

In this section...

“Introduction” on page 3-17
 “Common Headings” on page 3-18
 “Function-Specific Headings” on page 3-18

Introduction

Iterative display is a table of statistics describing the calculations in each iteration of a solver. The statistics depend on both the solver and the solver algorithm. For more information about iterations, see “Iterations and Function Counts” on page 3-11. The table appears in the MATLAB Command Window when you run solvers with appropriate options.

Obtain iterative display by using `optimoptions` to create options with the `Display` option set to `'iter'` or `'iter-detailed'`. For example:


```
options = optimoptions(@fminunc,'Display','iter','Algorithm','quasi-newton')
[x fval exitflag output] = fminunc(@sin,0,options);
```

Iteration	Func-count	f(x)	Step-size	First-order optimality
0	2	0		1
1	4	-0.841471	1	0.54
2	8	-1	0.484797	0.000993
3	10	-1	1	5.62e-005
4	12	-1	1	0

Local minimum found.

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

You can also obtain iterative display by using the Optimization app. Select **Display to command window > Level of display > iterative** or **iterative with detailed message**.



Iterative display is available for all solvers except:

- `linprog` medium-scale active-set algorithm
- `lsqlin`
- `lsqnonneg`
- `quadprog` trust-region-reflective and active-set algorithms

Common Headings

The following table lists some common headings of iterative display.

Heading	Information Displayed
$f(x)$	Current objective function value
First-order optimality	First-order optimality measure (see “First-Order Optimality Measure” on page 3-12)
Func-count or F-count	Number of function evaluations; see “Iterations and Function Counts” on page 3-11
Iteration or Iter	Iteration number; see “Iterations and Function Counts” on page 3-11
Norm of step	Size of the current step (size is the Euclidean norm, or 2-norm)

Function-Specific Headings

The following sections describe headings of iterative display whose meaning is specific to the optimization function you are using:

- “`fgoalattain`, `fmincon`, `fminimax`, and `fseminf`” on page 3-19
- “`fminbnd` and `fzero`” on page 3-21

- “fminsearch” on page 3-21
- “fminunc” on page 3-22
- “fsolve” on page 3-22
- “intlinprog” on page 3-23
- “linprog” on page 3-23
- “lsqnonlin and lsqcurvefit” on page 3-24
- “quadprog” on page 3-24

fgoalattain, fmincon, fminimax, and fsemif

The following table describes the headings specific to `fgoalattain`, `fmincon`, `fminimax`, and `fsemif`.

fgoalattain, fmincon, fminimax, or fsemif Heading	Information Displayed
Attainment factor	Value of the attainment factor for <code>fgoalattain</code> .
CG-iterations	Number of conjugate gradient iterations taken in the current iteration (see “Preconditioned Conjugate Gradient Method” on page 6-29).
Directional derivative	Gradient of the objective function along the search direction.
Feasibility	Maximum constraint violation, where satisfied inequality constraints count as 0.
Line search steplength	Multiplicative factor that scales the search direction (see Equation 6-46).
Max constraint	Maximum violation among all constraints, both internally constructed and user-provided; can be negative when no constraint is binding.
Objective value	Objective function value of the nonlinear programming reformulation of the minimax problem for <code>fminimax</code> .

fgoalattain, fmincon, fminimax, or fseminf Heading	Information Displayed
Procedure	<p>Hessian update procedures:</p> <ul style="list-style-type: none"> • Infeasible start point • Hessian not updated • Hessian modified • Hessian modified twice <p>For more information, see “Updating the Hessian Matrix” on page 6-35.</p> <p>QP subproblem procedures:</p> <ul style="list-style-type: none"> • dependent — There are dependent (redundant) equality constraints that the solver detected and removed. • Infeasible — The QP subproblem with linearized constraints is infeasible. • Overly constrained — The QP subproblem with linearized constraints is infeasible. • Unbounded — The QP subproblem is feasible with large negative curvature. • Ill-posed — The QP subproblem search direction is too small. • Unreliable — The QP subproblem seems to be ill-conditioned.
Steplength	Multiplicative factor that scales the search direction (see Equation 6-46).
Trust-region radius	Current trust-region radius.

fminbnd and fzero

The following table describes the headings specific to `fminbnd` and `fzero`.

fminbnd or fzero Heading	Information Displayed
Procedure	Procedures for <code>fminbnd</code> : <ul style="list-style-type: none"> • <code>initial</code> • <code>golden</code> (golden section search) • <code>parabolic</code> (parabolic interpolation) Procedures for <code>fzero</code> : <ul style="list-style-type: none"> • <code>initial</code> (initial point) • <code>search</code> (search for an interval containing a zero) • <code>bisection</code> • <code>interpolation</code> (linear interpolation or inverse quadratic interpolation)
<code>x</code>	Current point for the algorithm

fminsearch

The following table describes the headings specific to `fminsearch`.

fminsearch Heading	Information Displayed
<code>min f(x)</code>	Minimum function value in the current simplex.
Procedure	Simplex procedure at the current iteration. Procedures include: <ul style="list-style-type: none"> • <code>initial simplex</code> • <code>expand</code> • <code>reflect</code> • <code>shrink</code>

fminsearch Heading	Information Displayed
	<ul style="list-style-type: none"> • contract inside • contract outside <p>For details, see “fminsearch Algorithm” on page 6-14.</p>

fminunc

The following table describes the headings specific to fminunc.

fminunc Heading	Information Displayed
CG-iterations	Number of conjugate gradient iterations taken in the current iteration (see “Preconditioned Conjugate Gradient Method” on page 6-29)
Step-size	Multiplicative factor that scales the search direction (see Equation 6-12)

The fminunc medium-scale algorithm can issue a `skipped update` message to the right of the `First-order optimality` column. This message means that fminunc did not update its Hessian estimate, because the resulting matrix would not have been positive definite. The message usually indicates that the objective function is not smooth at the current point.

fsolve

The following table describes the headings specific to fsolve.

fsolve Heading	Information Displayed
Directional derivative	Gradient of the function along the search direction
Lambda	λ_k value defined in “Levenberg-Marquardt Method” on page 6-205

fsolve Heading	Information Displayed
Residual	Residual (sum of squares) of the function
Trust-region radius	Current trust-region radius (change in the norm of the trust-region radius)

intlinprog

The following table describes the headings specific to `intlinprog`.

linprog Heading	Information Displayed
nodes explored	Cumulative number of explored nodes.
total time (s)	Time in seconds since <code>intlinprog</code> started.
num int solution	Number of integer feasible points found.
integer fval	Objective function value of the best integer feasible point found. This is an upper bound for the final objective function value.
relative gap (%)	$\frac{100(b - a)}{ b + 1},$ <p>where</p> <ul style="list-style-type: none"> • b is the objective function value of the best integer feasible point. • a is the best lower bound on the objective function value.

linprog

The following table describes the headings specific to `linprog`.

linprog Heading	Information Displayed
Dual Infeas $A' * y + z - w - f$	Dual infeasibility.
Duality Gap $x' * z + s' * w$	Duality gap (see “Interior-Point Linear Programming” on page 6-106) between the primal objective and the dual objective. s and w appear only in this equation if there are finite upper bounds.
Objective $f' * x$	Current objective value.
Primal Infeas $A * x - b$	Primal infeasibility.
Total Rel Error	Total relative error, described at the end of “Main Algorithm” on page 6-106.

lsqnonlin and lsqcurvefit

The following table describes the headings specific to `lsqnonlin` and `lsqcurvefit`.

lsqnonlin or lsqcurvefit Heading	Information Displayed
Directional derivative	Gradient of the function along the search direction
Lambda	λ_k value defined in “Levenberg-Marquardt Method” on page 6-205
Resnorm	Value of the squared 2-norm of the residual at x
Residual	Residual vector of the function

quadprog

The following table describes the headings specific to `quadprog`.

quadprog Heading	Information Displayed
Feasibility	Maximum constraint violation, where satisfied inequality constraints count as 0.
Total relative error	Total relative error is a measure of infeasibility, as defined in "Total Relative Error" on page 6-125

Output Structures

An *output structure* contains information on a solver's result. All solvers can return an output structure. To obtain an output structure, invoke the solver with the output structure in the calling syntax. For example, to get an output structure from `lsqnonlin`, use the syntax

```
[x,resnorm,residual,exitflag,output] = lsqnonlin(...)
```

You can also obtain an output structure by running a problem using the Optimization app. All results exported from Optimization app contain an output structure.

The contents of the output structure are listed in each solver's reference pages. For example, the output structure returned by `lsqnonlin` contains `firstorderopt`, `iterations`, `funcCount`, `cgiterations`, `stepsize`, `algorithm`, and `message`. To access, for example, the message, enter `output.message`.

Optimization app exports results in a structure. The results structure contains the output structure. To access, for example, the number of iterations, use the syntax `optimresults.output.iterations`.

You can also see the contents of an output structure by double-clicking the output structure in the MATLAB Workspace pane.

Lagrange Multiplier Structures

Constrained optimization involves a set of Lagrange multipliers, as described in “First-Order Optimality Measure” on page 3-12. Solvers return estimated Lagrange multipliers in a structure. The structure is called `lambda`, since the conventional symbol for Lagrange multipliers is the Greek letter lambda (λ). The structure separates the multipliers into the following types, called fields:

- `lower`, associated with lower bounds
- `upper`, associated with upper bounds
- `eqlin`, associated with linear equalities
- `ineqlin`, associated with linear inequalities
- `eqnonlin`, associated with nonlinear equalities
- `ineqnonlin`, associated with nonlinear inequalities

To access, for example, the nonlinear inequality field of a Lagrange multiplier structure, enter `lambda.ineqnonlin`. To access the third element of the Lagrange multiplier associated with lower bounds, enter `lambda.lower(3)`.

The content of the Lagrange multiplier structure depends on the solver. For example, linear programming has no nonlinearities, so it does not have `eqnonlin` or `ineqnonlin` fields. Each applicable solver’s function reference pages contains a description of its Lagrange multiplier structure under the heading “Outputs.”

Hessian

In this section...

“fminunc Hessian” on page 3-28

“fmincon Hessian” on page 3-29

fminunc Hessian

The Hessian for an unconstrained problem is the matrix of second derivatives of the objective function f :

$$\text{Hessian } H_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}.$$

- **Quasi-Newton Algorithm** — `fminunc` returns an estimated Hessian matrix at the solution. It computes the estimate by finite differences.
- **Trust-Region Algorithm** — `fminunc` returns a Hessian matrix at the next-to-last iterate.
 - If you supply a Hessian in the objective function, `fminunc` returns this Hessian.
 - If you supply a `HessMult` function, `fminunc` returns the `Hinfo` matrix from the `HessMult` function. For more information, see “trust-region Algorithm Only” on page 10-103.
 - Otherwise, `fminunc` returns an approximation from a sparse finite difference algorithm on the gradients.

This Hessian is accurate for the next-to-last iterate. However, the next-to-last iterate might not be close to the final point.

The reason the trust-region algorithm returns the Hessian at the next-to-last point is for efficiency. `fminunc` uses the Hessian internally to compute its next step. When `fminunc` reaches a stopping condition, it does not need to compute the next step, so does not compute the Hessian.

fmincon Hessian

The Hessian for a constrained problem is the Hessian of the Lagrangian. For an objective function f , nonlinear inequality constraint vector c , and nonlinear equality constraint vector ceq , the Lagrangian is

$$L = f + \sum_i \lambda_i c_i + \sum_j \lambda_j ceq_j.$$

The λ_i are Lagrange multipliers; see “First-Order Optimality Measure” on page 3-12 and “Lagrange Multiplier Structures” on page 3-27. The Hessian of the Lagrangian is

$$H = \nabla^2 L = \nabla^2 f + \sum_i \lambda_i \nabla^2 c_i + \sum_j \lambda_j \nabla^2 ceq_j.$$

`fmincon` has four algorithms, with several options for Hessians, as described in “fmincon Trust Region Reflective Algorithm” on page 6-26, “fmincon Active Set Algorithm” on page 6-32, and “fmincon Interior Point Algorithm” on page 6-43. `fmincon` returns the following for the Hessian:

- **active-set or sqp Algorithm** — `fmincon` returns the Hessian approximation it computes at the next-to-last iterate. `fmincon` computes a quasi-Newton approximation of the Hessian matrix at the solution in the course of its iterations. This approximation does not, in general, match the true Hessian in every component, but only in certain subspaces. Therefore the Hessian that `fmincon` returns can be inaccurate. For more details of the active-set calculation, see “SQP Implementation” on page 6-35.
- **trust-region-reflective Algorithm** — `fmincon` returns the Hessian it computes at the next-to-last iterate.
 - If you supply a Hessian in the objective function, `fmincon` returns this Hessian.
 - If you supply a `HessMult` function, `fmincon` returns the `Hinfo` matrix from the `HessMult` function. For more information, see “Trust-Region-Reflective Algorithm” on page 10-57.
 - Otherwise, `fmincon` returns an approximation from a sparse finite difference algorithm on the gradients.

This Hessian is accurate for the next-to-last iterate. However, the next-to-last iterate might not be close to the final point.

The reason the `trust-region-reflective` algorithm returns the Hessian at the next-to-last point is for efficiency. `fmincon` uses the Hessian internally to compute its next step. When `fmincon` reaches a stopping condition, it does not need to compute the next step, so does not compute the Hessian.

- **interior-point Algorithm**

- If the Hessian option is `lbfgs` or `fin-diff-grads`, or if you supply a Hessian multiply function (`HessMult`), `fmincon` returns `[]` for the Hessian.
- If the Hessian option is `bfgs` (the default), `fmincon` returns a quasi-Newton approximation to the Hessian at the final point. This Hessian can be inaccurate, as in the `active-set` or `sqp` algorithm Hessian.
- If the Hessian option is `user-supplied`, `fmincon` returns the user-supplied Hessian at the final point.

Concepts

- “Including Derivatives” on page 2-22
- “Hessian” on page 10-49

Plot Functions

In this section...

“Plot an Optimization During Execution” on page 3-31

“Using a Plot Function” on page 3-31

Plot an Optimization During Execution

You can plot various measures of progress during the execution of a solver. Set the `PlotFcns` name-value pair in `optimoptions`, and specify one or more plotting functions for the solver to call at each iteration. Pass a function handle or cell array of function handles.

There are a variety of predefined plot functions available. See:

- The `PlotFcns` option description in the solver function reference page
- **Optimization app > Options > Plot functions**

You can also use a custom-written plot function. Write a function file using the same structure as an output function. For more information on this structure, see “Output Function” on page 9-21.

Using a Plot Function

This example shows how to use a plot function to view the progress of the `fmincon` interior-point algorithm. The problem is taken from the Getting Started “Solve a Constrained Nonlinear Problem” on page 1-3. The first part of the example shows how to run the optimization using the Optimization app. The second part shows how to run the optimization from the command line.

Running the Optimization Using the Optimization App

- 1 Write the nonlinear objective and constraint functions, including the derivatives:

```
function [f g H] = rosenboth(x)
% ROSENBOTh returns both the value y of Rosenbrock's function
% and also the value g of its gradient and H the Hessian.
```

```
f = 100*(x(2) - x(1)^2)^2 + (1-x(1))^2;

if nargout > 1
    g = [-400*(x(2)-x(1)^2)*x(1)-2*(1-x(1));
        200*(x(2)-x(1)^2)];

    if nargout > 2
        H = [1200*x(1)^2-400*x(2)+2, -400*x(1);
            -400*x(1), 200];
    end
end
```

Save this file as `rosenboth.m`.

```
function [c,ceq,gc,gceq] = unitdisk2(x)
% UNITDISK2 returns the value of the constraint
% function for the disk of radius 1 centered at
% [0 0]. It also returns the gradient.
```

```
c = x(1)^2 + x(2)^2 - 1;
ceq = [ ];
```

```
if nargout > 2
    gc = [2*x(1);2*x(2)];
    gceq = [ ];
end
```

Save this file as `unitdisk2.m`.

2 Start the Optimization app by entering `optimtool` at the command line.

3 Set up the optimization:

- Choose the `fmincon` solver.
- Choose the Interior point algorithm.
- Set the objective function to `@rosenboth`.
- Choose `Gradient` supplied for the objective function derivative.
- Set the start point to `[0 0]`.

- Set the nonlinear constraint function to @unitdisk2.
- Choose Gradient supplied for the nonlinear constraint derivatives.

Your **Problem Setup and Results** panel should match the following figure.

Problem Setup and Results

Solver:

Algorithm:

Problem

Objective function:

Derivatives:

Start point:

Constraints:

Linear inequalities: A: b:

Linear equalities: Aeq: beq:

Bounds: Lower: Upper:

Nonlinear constraint function:

Derivatives:

- 4** Choose three plot functions in the **Options** pane: **Current point**, **Function value**, and **First order optimality**.

Plot functions

Current point Function count Function value

Max constraint Current step First order optimality

Custom function:

- 5 Click the **Start** button under **Run solver and view results**.
- 6 The output appears as follows in the Optimization app.

Run solver and view results

Start Pause Stop

Current iteration: 24 Clear Results

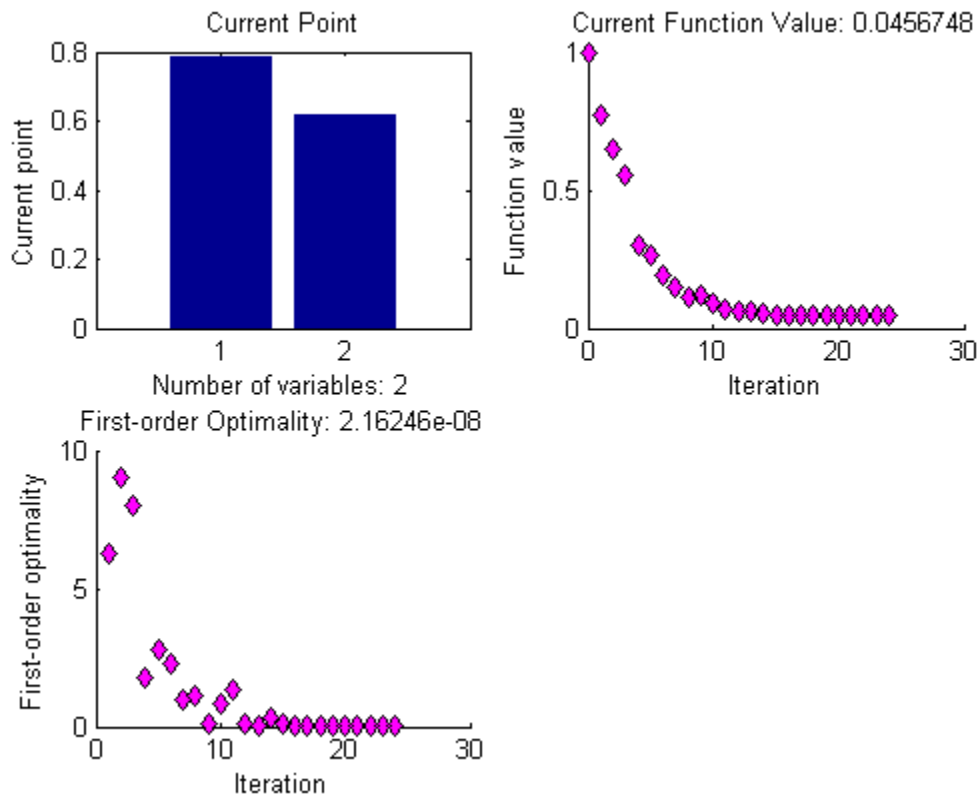
Optimization running.
Objective function value: 0.0456748247578765
Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.

Final point:

1 ▲	2
0.786	0.618

In addition, the following three plots appear in a separate window.



- The “Current Point” plot graphically shows the minimizer $[0.786, 0.618]$, which is reported as the **Final point** in the **Run solver and view results** pane. This plot updates at each iteration, showing the intermediate iterates.
- The “Current Function Value” plot shows the objective function value at all iterations. This graph is nearly monotone, showing `fmincon` reduces the objective function at almost every iteration.
- The “First-order Optimality” plot shows the first-order optimality measure at all iterations.

Running the Optimization from the Command Line

1 Write the nonlinear objective and constraint functions, including the derivatives, as shown in “Running the Optimization Using the Optimization App” on page 3-31.

2 Create an options structure that includes calling the three plot functions:

```
options = optimoptions(@fmincon,'Algorithm','interior-point',...  
    'GradObj','on','GradConstr','on','PlotFcns',{@optimplotx,...  
    @optimplotfval,@optimplotfirstorderopt});
```

3 Call `fmincon`:

```
x = fmincon(@rosenboth,[0 0],[[],[],[],[],[],[],[],...  
    @unitdisk2,options)
```

4 `fmincon` gives the following output:

```
Local minimum found that satisfies the constraints.
```

```
Optimization completed because the objective function is  
non-decreasing in feasible directions, to within the default  
value of the function tolerance, and constraints are satisfied  
to within the default value of the constraint tolerance.
```

```
x =  
    0.7864    0.6177
```

`fmincon` also displays the three plot functions, shown at the end of “Running the Optimization Using the Optimization App” on page 3-31.

Output Functions

In this section...
“What Is an Output Function?” on page 3-37
“Example: Using Output Functions” on page 3-37

What Is an Output Function?

For some problems, you might want output from an optimization algorithm at each iteration. For example, you might want to find the sequence of points that the algorithm computes and plot those points. To do this, create an output function that the optimization function calls at each iteration. See “Output Function” on page 9-21 for details and syntax.

Generally, the solvers that can employ an output function are the ones that can take nonlinear functions as inputs. You can determine which solvers can have an output function by looking in the Options section of function reference pages, or by checking whether the **Output function** option is available in the Optimization app for a solver.

Example: Using Output Functions

- “What the Example Contains” on page 3-37
- “Writing the Output Function” on page 3-38
- “Writing the Example Function File” on page 3-39
- “Running the Example” on page 3-41

What the Example Contains

The following example continues the one in “Nonlinear Inequality Constraints” on page 6-52, which calls the function `fmincon` at the command line to solve a nonlinear, constrained optimization problem. The example in this section uses a function file to call `fmincon`. The file also contains all the functions needed for the example, including:

- The objective function

- The constraint function
- An output function that records the history of points computed by the algorithm for `fmincon`. At each iteration of the algorithm for `fmincon`, the output function:
 - Plots the current point computed by the algorithm.
 - Stores the point and its corresponding objective function value in a variable called `history`, and stores the current search direction in a variable called `searchdir`. The search direction is a vector that points in the direction from the current point to the next one.

The code for the file is here: “Writing the Example Function File” on page 3-39.

Writing the Output Function

You specify the output function in options, such as

```
options = optimoptions(@fmincon, 'OutputFcn', @outfun)
```

where `outfun` is the name of the output function. When you call an optimization function with `options` as an input, the optimization function calls `outfun` at each iteration of its algorithm.

In general, `outfun` can be any MATLAB function, but in this example, it is a nested function of the function file described in “Writing the Example Function File” on page 3-39. The following code defines the output function:

```
function stop = outfun(x,optimValues,state)
stop = false;

    switch state
        case 'init'
            hold on
        case 'iter'
            % Concatenate current point and objective function
            % value with history. x must be a row vector.
            history.fval = [history.fval; optimValues.fval];
            history.x = [history.x; x];
            % Concatenate current search direction with
            % searchdir.
```

```

        searchdir = [searchdir;...
                    optimValues.searchdirection'];
        plot(x(1),x(2),'o');
        % Label points with iteration number.
        % Add .15 to x(1) to separate label from plotted 'o'
        text(x(1)+.15,x(2),num2str(optimValues.iteration));
    case 'done'
        hold off
    otherwise
    end
end
end

```

See “Using Handles to Store Function Parameters” in the MATLAB Programming Fundamentals documentation for more information about nested functions.

The arguments that the optimization function passes to `outfun` are:

- `x` — The point computed by the algorithm at the current iteration
- `optimValues` — Structure containing data from the current iteration
 - The example uses the following fields of `optimValues`:
 - `optimValues.iteration` — Number of the current iteration
 - `optimValues.fval` — Current objective function value
 - `optimValues.searchdirection` — Current search direction
- `state` — The current state of the algorithm ('init', 'interrupt', 'iter', or 'done')

For more information about these arguments, see “Output Function” on page 9-21.

Writing the Example Function File

To create the function file for this example:

- 1** Open a new file in the MATLAB Editor.
- 2** Copy and paste the following code into the file:

```
function [history,searchdir] = runfmincon

% Set up shared variables with OUTFUN
history.x = [];
history.fval = [];
searchdir = [];

% call optimization
x0 = [-1 1];
options = optimoptions(@fmincon,'OutputFcn',@outfun,...
    'Display','iter','Algorithm','active-set');
xsol = fmincon(@objfun,x0,[],[],[],[],[],[],@confun,options);

function stop = outfun(x,optimValues,state)
    stop = false;

    switch state
        case 'init'
            hold on
        case 'iter'
            % Concatenate current point and objective function
            % value with history. x must be a row vector.
            history.fval = [history.fval; optimValues.fval];
            history.x = [history.x; x];
            % Concatenate current search direction with
            % searchdir.
            searchdir = [searchdir;...
                optimValues.searchdirection'];
            plot(x(1),x(2),'o');
            % Label points with iteration number and add title.
            % Add .15 to x(1) to separate label from plotted 'o'
            text(x(1)+.15,x(2),...
                num2str(optimValues.iteration));
            title('Sequence of Points Computed by fmincon');
        case 'done'
            hold off
        otherwise
    end
end
```

```

function f = objfun(x)
    f = exp(x(1))*(4*x(1)^2 + 2*x(2)^2 + 4*x(1)*x(2) + ...
           2*x(2) + 1);
end

function [c, ceq] = confun(x)
    % Nonlinear inequality constraints
    c = [1.5 + x(1)*x(2) - x(1) - x(2);
         -x(1)*x(2) - 10];
    % Nonlinear equality constraints
    ceq = [];
end
end

```

3 Save the file as `runfmincon.m` in a folder on the MATLAB path.

Running the Example

To run the example, enter:

```
[history searchdir] = runfmincon;
```

This displays the following iterative output in the Command Window.

Iter	F-count	f(x)	Max constraint	Line search steplength	Directional derivative	First-order optimality	Procedure
0	3	1.8394	0.5				Infeasible
1	6	1.85127	-0.09197	1	0.109	0.778	start point
2	9	0.300167	9.33	1	-0.117	0.313	Hessian modified
3	12	0.529835	0.9209	1	0.12	0.232	twice
4	16	0.186965	-1.517	0.5	-0.224	0.13	
5	19	0.0729085	0.3313	1	-0.121	0.054	
6	22	0.0353323	-0.03303	1	-0.0542	0.0271	
7	25	0.0235566	0.003184	1	-0.0271	0.00587	
8	28	0.0235504	9.032e-008	1	-0.0146	8.51e-007	

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.

```
Active inequalities (to within options.TolCon = 1e-006):
```

```
  lower      upper      ineqlin      ineqnonlin
                                     1
                                     2
```

The output `history` is a structure that contains two fields:

```
history =
        x: [9x2 double]
       fval: [9x1 double]
```

The `fval` field contains the objective function values corresponding to the sequence of points computed by `fmincon`:

```
history.fval
```

```
ans =
    1.8394
    1.8513
    0.3002
    0.5298
    0.1870
    0.0729
    0.0353
    0.0236
    0.0236
```

These are the same values displayed in the iterative output in the column with header `f(x)`.

The `x` field of `history` contains the sequence of points computed by the algorithm:

```
history.x
ans =
    -1.0000    1.0000
    -1.3679    1.2500
```




```

-5.5708    3.4699
-4.8000    2.2752
-6.7054    1.2618
-8.0679    1.0186
-9.0230    1.0532
-9.5471    1.0471
-9.5474    1.0474

```

This example displays a plot of this sequence of points, in which each point is labeled by its iteration number.



The optimal point occurs at the eighth iteration. Note that the last two points in the sequence are so close that they overlap.

The second output argument, `searchdir`, contains the search directions for `fmincon` at each iteration. The search direction is a vector pointing from the point computed at the current iteration to the point computed at the next iteration:

```

searchdir =

    -0.3679    0.2500

```

-4.2029	2.2199
0.7708	-1.1947
-3.8108	-2.0268
-1.3625	-0.2432
-0.9552	0.0346
-0.5241	-0.0061
-0.0003	0.0003

Steps to Take After Running a Solver

- “Overview of Next Steps” on page 4-2
- “When the Solver Fails” on page 4-3
- “When the Solver Might Have Succeeded” on page 4-15
- “When the Solver Succeeds” on page 4-23
- “Optimizing a Simulation or Ordinary Differential Equation” on page 4-33

Overview of Next Steps

This topic addresses questions you might have after running a solver. The questions include:

- Is the answer reliable?
- What can you do if the solver fails?
- Is the minimum smaller than all other minima, or only smaller than nearby minima? (“Local vs. Global Optima” on page 4-27)
- What can you do if the solver takes too long?

The list of questions is not exhaustive. It covers common or basic problems.

You can access relevant answers from many solvers’ default exit message. The first line of the exit message contains a link to a brief description of the result. This description contains a link leading to documentation.

When the Solver Fails

In this section...
“Too Many Iterations or Function Evaluations” on page 4-3
“No Feasible Point” on page 4-7
“Problem Unbounded” on page 4-9
“fsolve Could Not Solve Equation” on page 4-10
“Solver Takes Too Long” on page 4-10

Too Many Iterations or Function Evaluations

The solver stopped because it reached a limit on the number of iterations or function evaluations before it minimized the objective to the requested tolerance. To proceed, try one or more of the following.

- “1. Enable Iterative Display” on page 4-3
- “2. Relax Tolerances” on page 4-4
- “3. Start the Solver From Different Points” on page 4-5
- “4. Check Objective and Constraint Function Definitions” on page 4-5
- “5. Center and Scale Your Problem” on page 4-5
- “6. Provide Gradient or Jacobian” on page 4-6
- “7. Provide Hessian” on page 4-7

1. Enable Iterative Display

Set the `Display` option to `'iter'`. This setting shows the results of the solver iterations.

To enable iterative display:

- Using the Optimization app, choose **Level of display** to be **iterative or iterative with detailed message**.
- At the MATLAB command line, enter

```
options = optimoptions('solvername','Display','iter');
```

Call the solver using the options structure.

For an example of iterative display, see “Interpreting the Result” on page 1-11.

What to Look For in Iterative Display.

- See if the objective function (Fval or $f(x)$ or Resnorm) decreases. Decrease indicates progress.
- Examine constraint violation (Max constraint) to ensure that it decreases towards 0. Decrease indicates progress.
- See if the first-order optimality decreases towards 0. Decrease indicates progress.
- See if the Trust-region radius decreases to a small value. This decrease indicates that the objective might not be smooth.

What to Do.

- If the solver seemed to progress:
 - 1 Set MaxIter and/or MaxFunEvals to values larger than the defaults. You can see the default values in the Optimization app, or in the Options table in the solver’s function reference pages.
 - 2 Start the solver from its last calculated point.
- If the solver is not progressing, try the other listed suggestions.

2. Relax Tolerances

If TolX or TolFun, for example, are too small, the solver might not recognize when it has reached a minimum; it can make futile iterations indefinitely.

To change tolerances using the Optimization app, use the **Stopping criteria** list at the top of the **Options** pane.

To change tolerances at the command line, use optimoptions as described in “Set and Change Options” on page 2-59.

The DiffMaxChange and DiffMinChange options can affect a solver’s progress. These options control the step size in finite differencing for derivative estimation.

3. Start the Solver From Different Points

See Change the Initial Point.

4. Check Objective and Constraint Function Definitions

For example, check that your objective and nonlinear constraint functions return the correct values at some points. See Check your Objective and Constraint Functions. Check that an infeasible point does not cause an error in your functions; see “Iterations Can Violate Constraints” on page 2-34.

5. Center and Scale Your Problem

Solvers run more reliably when each coordinate has about the same effect on the objective and constraint functions. Multiply your coordinate directions with appropriate scalars to equalize the effect of each coordinate. Add appropriate values to certain coordinates to equalize their size.

Example: Centering and Scaling. Consider minimizing $1e6*x(1)^2 + 1e-6*x(2)^2$:

```
f = @(x) 10^6*x(1)^2 + 10^-6*x(2)^2;
```

Minimize f using the medium-scale `fminunc` algorithm:

```
opts = optimoptions('fminunc','Display','none','Algorithm','quasi-newton');
x = fminunc(f,[0.5;0.5],opts)
```

```
x =
      0
 0.5000
```

The result is incorrect; poor scaling interfered with obtaining a good solution.

Scale the problem. Set

```
D = diag([1e-3,1e3]);
fr = @(y) f(D*y);
y = fminunc(fr, [0.5;0.5], opts)
```

```
y =
      0
```

```
0 % the correct answer
```

Similarly, poor centering can interfere with a solution.

```
fc = @(z)fr([z(1)-1e6;z(2)+1e6]); % poor centering
z = fminunc(fc,[.5 .5],opts)
```

```
z =
  1.0e+005 *
  10.0000 -10.0000 % looks good, but...
```

```
z - [1e6 -1e6] % checking how close z is to 1e6
```

```
ans =
-0.0071    0.0078 % reveals a distance
```

```
fcc = @(w)fc([w(1)+1e6;w(2)-1e6]); % centered
```

```
w = fminunc(fcc,[.5 .5],opts)
```

```
w =
  0    0 % the correct answer
```

6. Provide Gradient or Jacobian

If you do not provide gradients or Jacobians, solvers estimate gradients and Jacobians by finite differences. Therefore, providing these derivatives can save computational time, and can lead to increased accuracy.

For constrained problems, providing a gradient has another advantage. A solver can reach a point x such that x is feasible, but finite differences around x always lead to an infeasible point. In this case, a solver can fail or halt prematurely. Providing a gradient allows a solver to proceed.

Provide gradients or Jacobians in the files for your objective function and nonlinear constraint functions. For details of the syntax, see “Writing Scalar Objective Functions” on page 2-19, “Writing Vector and Matrix Objective Functions” on page 2-27, and “Nonlinear Constraints” on page 2-37.

To check that your gradient or Jacobian function is correct, use the `DerivativeCheck` option, as described in “Checking Validity of Gradients or Jacobians” on page 2-69.

If you have a Symbolic Math Toolbox license, you can calculate gradients and Hessians programmatically. For an example, see “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85.

For examples using gradients and Jacobians, see “Minimization with Gradient and Hessian” on page 6-20, “Nonlinear Constraints with Gradients” on page 6-54, “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85, “Nonlinear Equations with Analytic Jacobian” on page 6-257, and “Nonlinear Equations with Jacobian” on page 6-263.

7. Provide Hessian

Solvers often run more reliably and with fewer iterations when you supply a Hessian.

The following solvers and algorithms accept Hessians:

- `fmincon interior-point`. Write the Hessian as a separate function. For an example, see “`fmincon Interior-Point Algorithm with Analytic Hessian`” on page 6-57.
- `fmincon trust-region-reflective`. Give the Hessian as the third output of the objective function. For an example, see “`Minimization with Dense Structured Hessian, Linear Equalities`” on page 6-80.
- `fminunc trust-region`. Give the Hessian as the third output of the objective function. For an example, see “`Minimization with Gradient and Hessian`” on page 6-20.

If you have a Symbolic Math Toolbox license, you can calculate gradients and Hessians programmatically. For an example, see “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85.

No Feasible Point

The solver was unable to find a point satisfying all constraints to within the `TolCon` constraint tolerance. To proceed, try one or more of the following.

- “1. Check Linear Constraints” on page 4-8
- “2. Check Nonlinear Constraints” on page 4-8

1. Check Linear Constraints

Try finding a point that satisfies the bounds and linear constraints by solving a linear programming problem.

- 1 Define a linear programming problem with an objective function that is always zero:

```
f = zeros(size(x0)); % assumes x0 is the initial point
```

- 2 Solve the linear programming problem to see if there is a feasible point:

```
xnew = linprog(f,A,b,Aeq,beq,lb,ub);
```

- 3 If there is a feasible point x_{new} , use x_{new} as the initial point and rerun your original problem.
- 4 If there is no feasible point, your problem is not well-formulated. Check the definitions of your bounds and linear constraints.

2. Check Nonlinear Constraints

After ensuring that your bounds and linear constraints are feasible (contain a point satisfying all constraints), check your nonlinear constraints.

- Set your objective function to zero:

```
@(x)0
```

Run your optimization with the zero objective. If you find a feasible point x_{new} , set $x_0 = x_{new}$ and rerun your original problem.

- If you do not find a feasible point using a zero objective function, use the zero objective function with several initial points.
 - If you find a feasible point x_{new} , set $x_0 = x_{new}$ and rerun your original problem.
 - If you do not find a feasible point, try relaxing the constraints, discussed next.

Try relaxing your nonlinear inequality constraints, then tightening them.

- 1 Change the nonlinear constraint function c to return $c - \Delta$, where Δ is a positive number. This change makes your nonlinear constraints easier to satisfy.
- 2 Look for a feasible point for the new constraint function, using either your original objective function or the zero objective function.
 - a If you find a feasible point,
 - a Reduce Δ
 - b Look for a feasible point for the new constraint function, starting at the previously found point.
 - b If you do not find a feasible point, try increasing Δ and looking again.

If you find no feasible point, your problem might be truly infeasible, meaning that no solution exists. Check all your constraint definitions again.

Problem Unbounded

The solver reached a point whose objective function was less than the objective limit tolerance.

- Your problem might be truly unbounded. In other words, there is a sequence of points x_i with

$$\lim f(x_i) = -\infty.$$

and such that all the x_i satisfy the problem constraints.

- Check that your problem is formulated correctly. Solvers try to minimize objective functions; if you want a maximum, change your objective function to its negative. For an example, see “Maximizing an Objective” on page 2-31.
- Try scaling or centering your problem. See Center and Scale Your Problem.
- Relax the objective limit tolerance by using `optimoptions` to reduce the value of the `ObjectiveLimit` tolerance.

fsolve Could Not Solve Equation

`fsolve` can fail to solve an equation for various reasons. Here are some suggestions for how to proceed:

- 1 Try Changing the Initial Point.** `fsolve` relies on an initial point. By giving it different initial points, you increase the chances of success.
- 2 Check the definition of the equation to make sure that it is smooth.** `fsolve` might fail to converge for equations with discontinuous gradients, such as absolute value. `fsolve` can fail to converge for functions with discontinuities.
- 3 Check that the equation is “square,”** meaning equal dimensions for input and output (has the same number of unknowns as values of the equation).
- 4 Change tolerances, especially TolFun and TolX.** If you attempt to get high accuracy by setting tolerances to very small values, `fsolve` can fail to converge. If you set tolerances that are too high, `fsolve` can fail to solve an equation accurately.
- 5 Check the problem definition.** Some problems have no real solution, such as $x^2 + 1 = 0$.

Solver Takes Too Long

Solvers can take excessive time for various reasons. To diagnose the reason, use one or more of the following techniques.

- 1 “Enable Iterative Display”** on page 4-11
- 2 “Enable FunValCheck”** on page 4-11
- 3 “Use Appropriate Tolerances”** on page 4-11
- 4 “Use a Plot Function”** on page 4-12
- 5 “Enable DerivativeCheck”** on page 4-12
- 6 “Use Inf Instead of a Large, Arbitrary Bound”** on page 4-12
- 7 “Use an Output Function”** on page 4-13

8 “Use a Sparse Solver or a Multiply Function” on page 4-13

9 “Use Parallel Computing” on page 4-14

Enable Iterative Display

Set the `Display` option to `'iter'`. This setting shows the results of the solver iterations.

To enable iterative display:

- Using the Optimization app, choose **Level of display** to be **iterative** or **iterative with detailed message**.
- At the MATLAB command line, enter

```
options = optimoptions('solvername','Display','iter');
```

Call the solver using the options structure.

For an example of iterative display, see “Interpreting the Result” on page 1-11. For more information, see “What to Look For in Iterative Display” on page 4-4.

Enable FunValCheck

Sometimes a solver fails because an objective function or nonlinear constraint function returns a complex value, infinity, or NaN. To halt solver iterations in these cases, enable the `FunValCheck` option.

- Using the Optimization app, check the box labeled **Error if user-supplied function returns Inf, NaN, or complex** in the **Function value check** pane.
- At the MATLAB command line, enter

```
options = optimoptions('solvername','FunValCheck','on');
```

Call the solver using the options structure.

Use Appropriate Tolerances

Solvers can fail to converge if tolerances are too small, especially `TolFun` and `TolX`.

To change tolerances using the Optimization app, use the **Stopping criteria** list at the top of the **Options** pane.

To change tolerances at the command line, use `optimoptions` as described in “Set and Change Options” on page 2-59.

Use a Plot Function

You can obtain more visual or detailed information about solver iterations using a plot function. For a list of the predefined plot functions, see **Options > Plot functions** in the Optimization app. The Options section of your solver’s function reference pages also lists the plot functions.

To use a plot function:

- Using the Optimization app, check the boxes next to each plot function you wish to use.
- At the MATLAB command line, enter

```
options =  
optimoptions('solvername', 'PlotFcns', {@plotfcn1, @plotfcn2, ...});
```

Call the solver using the options structure.

For an example of using a plot function, see “Using a Plot Function” on page 3-31.

Enable DerivativeCheck

If you have supplied derivatives (gradients or Jacobians) to your solver, the solver can fail to converge if the derivatives are inaccurate. For more information about using the `DerivativeCheck` option, see “Checking Validity of Gradients or Jacobians” on page 2-69.

Use Inf Instead of a Large, Arbitrary Bound

If you use a large, arbitrary bound (upper or lower), a solver can take excessive time, or even fail to converge. However, if you set `Inf` or `-Inf` as the bound, the solver can take less time, and might converge better.

Why? An interior-point algorithm can set an initial point to the midpoint of finite bounds. Or an interior-point algorithm can try to find a “central path” midway between finite bounds. Therefore, a large, arbitrary bound can resize those components inappropriately. In contrast, infinite bounds are ignored for these purposes.

Minor point: Some solvers use memory for each constraint, primarily via a constraint Hessian. Setting a bound to `Inf` or `-Inf` means there is no constraint, so there is less memory in use, because a constraint Hessian has lower dimension.

Use an Output Function

You can obtain detailed information about solver iterations using an output function. Solvers call output functions at each iteration. You write output functions using the syntax described in “Output Function” on page 9-21.

For an example of using an output function, see “Example: Using Output Functions” on page 3-37.

Use a Sparse Solver or a Multiply Function

Large problems can cause MATLAB to run out of memory or time. Here are some suggestions for using less memory:

- Use a large-scale algorithm if possible (see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12). These algorithms include `trust-region-reflective`, `interior-point`, the `fminunc` `trust-region` algorithm, the `fsolve` `trust-region-dogleg` algorithm, and the Levenberg-Marquardt algorithm. In contrast, the `active-set`, `quasi-newton`, and `sqp` algorithms are not large-scale.

Use sparse matrices for your linear constraints.

- Use a Jacobian multiply function or Hessian multiply function. For examples, see “Jacobian Multiply Function with Linear Least Squares” on page 6-224, “Quadratic Minimization with Dense, Structured Hessian” on page 6-138, and “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80.

Use Parallel Computing

If you have a Parallel Computing Toolbox license, your solver might run faster using parallel computing. For more information, see “Parallel Computing”.

When the Solver Might Have Succeeded

In this section...
“Final Point Equals Initial Point” on page 4-15
“Local Minimum Possible” on page 4-15

Final Point Equals Initial Point

The initial point seems to be a local minimum or solution because the first-order optimality measure is close to 0. You might be unhappy with this result, since the solver did not improve your initial point.

If you are unsure that the initial point is truly a local minimum, try:

- 1 Starting from different points — see [Change the Initial Point](#).
- 2 Checking that your objective and constraints are defined correctly (for example, do they return the correct values at some points?) — see [Check your Objective and Constraint Functions](#). Check that an infeasible point does not cause an error in your functions; see [“Iterations Can Violate Constraints”](#) on page 2-34.
- 3 Changing tolerances, such as TolFun, TolCon, and TolX — see [Use Appropriate Tolerances](#).
- 4 Scaling your problem so each coordinate has about the same effect — see [Rescale the Problem](#).
- 5 Providing gradient and Hessian information — see [Provide Analytic Gradients or Jacobian](#) and [Provide a Hessian](#).

Local Minimum Possible

The solver might have reached a local minimum, but cannot be certain because the first-order optimality measure is not less than the TolFun tolerance. (To learn more about first-order optimality measure, see [“First-Order Optimality Measure”](#) on page 3-12.) To see if the reported solution is reliable, consider the following suggestions.

- “1. Nonsmooth Functions” on page 4-16

- “2. Rerun Starting At Final Point” on page 4-16
- “3. Try a Different Algorithm” on page 4-17
- “4. Change Tolerances” on page 4-19
- “5. Rescale the Problem” on page 4-20
- “6. Check Nearby Points” on page 4-20
- “7. Change Finite Differencing Options” on page 4-20
- “8. Provide Analytic Gradients or Jacobian” on page 4-21
- “9. Provide a Hessian” on page 4-21

1. Nonsmooth Functions

If you try to minimize a nonsmooth function, or have nonsmooth constraints, “Local Minimum Possible” can be the best exit message. This is because the first-order optimality conditions do not apply at a nonsmooth point.

To satisfy yourself that the solution is adequate, try to Check Nearby Points.

2. Rerun Starting At Final Point

Restarting an optimization at the final point can lead to a solution with a better first-order optimality measure. A better (lower) first-order optimality measure gives you more reason to believe that the answer is reliable.

For example, consider the following minimization problem, taken from an example (echodemo symbolic_optim_demo):

```
options = optimoptions('fminunc','Algorithm','quasi-newton');
funh = @(x)log(1 + (x(1) - 4/3)^2 + 3*(x(2) - (x(1)^3 - x(1)))^2);
[xfinal fval exitflag] = fminunc(funh,[-1;2],options)
```

```
Local minimum possible.
```

```
fminunc stopped because it cannot decrease the
objective function along the current search direction.
```

```
xfinal =
    1.3333
    1.0370
```

```
fval =
```

```
8.5265e-014
```

```
exitflag =
    5
```

The exit flag value of 5 indicates that the first-order optimality measure was above the function tolerance. Run the minimization again starting from `xfinal`:

```
[xfinal2 fval2 exitflag2] = fminunc(funh,xfinal,options)
```

```
Local minimum found.
```

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```
xfinal2 =
    1.3333
    1.0370
```

```
fval2 =
    6.5281e-014
```

```
exitflag2 =
    1
```

The local minimum is “found,” not “possible,” and the exitflag is 1, not 5. The two solutions are virtually identical. Yet the second run has a more satisfactory exit message, since the first-order optimality measure was low enough: $7.5996e-007$, instead of $3.9674e-006$.

3. Try a Different Algorithm

Many solvers give you a choice of algorithm. Different algorithms can lead to the use of different stopping criteria.

For example, `Rerun Starting At Final Point` returns exitflag 5 from the first run. This run uses the medium-scale algorithm.

The trust-region algorithm requires a user-supplied gradient. `betopt.m` contains a calculation of the objective function and gradient:

```
function [f gradf] = betopt(x)

g = 1 + (x(1)-4/3)^2 + 3*(x(2) - (x(1)^3-x(1)))^2;
f = log(g);
gradf(1) = 2*(x(1)-4/3) + 6*(x(2) - (x(1)^3-x(1)))*(1-3*x(1)^2);
gradf(1) = gradf(1)/g;
gradf(2) = 6*(x(2) - (x(1)^3 -x(1)))/g;
```

Running the optimization using the trust-region algorithm results in a different exitflag:

```
options = optimoptions('fminunc','GradObj','on');
[xfinal3 fval3 exitflag3] = fminunc(@betopt,[-1;2],options)
```

Local minimum possible.

fminunc stopped because the final change in function value relative to its initial value is less than the default value of the function tolerance.

```
xfinal3 =
    1.3333
    1.0370

fval3 =
    7.6659e-012

exitflag3 =
    3
```

The exit condition is better than the medium-scale condition, though it is still not the best. Rerunning the algorithm from the final point produces a better point, with extremely small first-order optimality measure:

```
[xfinal4 fval4 eflag4 output4] = fminunc(@betopt,xfinal3,options)
```

Local minimum found.

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```

xfinal4 =
    1.3333
    1.0370

fval4 =
    0

eflag4 =
    1

output4 =
    iterations: 1
    funcCount: 2
    cgiterations: 1
    firstorderopt: 7.5211e-011
    algorithm: 'large-scale: trust-region Newton'
    message: [1x498 char]

```

4. Change Tolerances

Sometimes tightening or loosening tolerances leads to a more satisfactory result. For example, choose a smaller value of TolFun in the Try a Different Algorithm section:

```

options = optimoptions('fminunc','TolFun',1e-8,'GradObj','on'); % default=1e-6
[xfinal3 fval3 eflag3 output3] = fminunc(@betopt,[-1;2],options)

```

Local minimum found.

Optimization completed because the size of the gradient is less than the selected value of the function tolerance.

```

xfinal3 =
    1.3333
    1.0370

fval3 =
    0

```

```
eflag3 =  
    1  
  
output3 =  
    iterations: 15  
    funcCount: 16  
    cgiterations: 12  
    firstorderopt: 7.5497e-11  
    algorithm: 'large-scale: trust-region Newton'  
    message: [1x498 char]  
    constrviolation: []
```

fminunc took one more iteration than before, arriving at a better solution.

5. Rescale the Problem

Try to have each coordinate give about the same effect on the objective and constraint functions by scaling and centering. For examples, see [Center and Scale Your Problem](#).

6. Check Nearby Points

Evaluate your objective function and constraints, if they exist, at points near the final point. If the final point is a local minimum, nearby feasible points have larger objective function values. See [Check Nearby Points](#) for an example.

If you have a Global Optimization Toolbox license, try running the `patternsearch` solver from the final point. `patternsearch` examines nearby points, and accepts all types of constraints.

7. Change Finite Differencing Options

Central finite differences are more time-consuming to evaluate, but are much more accurate. Use central differences when your problem can have high curvature.

To choose central differences at the command line, use `optimoptions` to set `'FinDiffType'` to `'central'`, instead of the default `'forward'`.

To choose central differences in the Optimization app, set **Options > Approximated derivatives > Type** to be central differences.

8. Provide Analytic Gradients or Jacobian

If you do not provide gradients or Jacobians, solvers estimate gradients and Jacobians by finite differences. Therefore, providing these derivatives can save computational time, and can lead to increased accuracy.

For constrained problems, providing a gradient has another advantage. A solver can reach a point x such that x is feasible, but finite differences around x always lead to an infeasible point. In this case, a solver can fail or halt prematurely. Providing a gradient allows a solver to proceed.

Provide gradients or Jacobians in the files for your objective function and nonlinear constraint functions. For details of the syntax, see “Writing Scalar Objective Functions” on page 2-19, “Writing Vector and Matrix Objective Functions” on page 2-27, and “Nonlinear Constraints” on page 2-37.

To check that your gradient or Jacobian function is correct, use the `DerivativeCheck` option, as described in “Checking Validity of Gradients or Jacobians” on page 2-69.

If you have a Symbolic Math Toolbox license, you can calculate gradients and Hessians programmatically. For an example, see “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85.

For examples using gradients and Jacobians, see “Minimization with Gradient and Hessian” on page 6-20, “Nonlinear Constraints with Gradients” on page 6-54, “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85, “Nonlinear Equations with Analytic Jacobian” on page 6-257, and “Nonlinear Equations with Jacobian” on page 6-263.

9. Provide a Hessian

Solvers often run more reliably and with fewer iterations when you supply a Hessian.

The following solvers and algorithms accept Hessians:

- `fmincon interior-point`. Write the Hessian as a separate function. For an example, see “`fmincon Interior-Point Algorithm with Analytic Hessian`” on page 6-57.
- `fmincon trust-region-reflective`. Give the Hessian as the third output of the objective function. For an example, see “`Minimization with Dense Structured Hessian, Linear Equalities`” on page 6-80.
- `fminunc trust-region`. Give the Hessian as the third output of the objective function. For an example, see “`Minimization with Gradient and Hessian`” on page 6-20.

If you have a Symbolic Math Toolbox license, you can calculate gradients and Hessians programmatically. For an example, see “`Symbolic Math Toolbox Calculates Gradients and Hessians`” on page 6-85.

The example in “`Symbolic Math Toolbox Calculates Gradients and Hessians`” on page 6-85 shows `fmincon` taking 77 iterations without a Hessian, but only 19 iterations with a Hessian.

When the Solver Succeeds

In this section...

“What Can Be Wrong If The Solver Succeeds?” on page 4-23

“1. Change the Initial Point” on page 4-24

“2. Check Nearby Points” on page 4-25

“3. Check your Objective and Constraint Functions” on page 4-26

“Local vs. Global Optima” on page 4-27

What Can Be Wrong If The Solver Succeeds?

A solver can report that a minimization succeeded, and yet the reported solution can be incorrect. For a rather trivial example, consider minimizing the function $f(x) = x^3$ for x between -2 and 2 , starting from the point $1/3$:

```
options = optimoptions('fmincon','Algorithm','active-set');
ffun = @(x)x^3;
xfinal = fmincon(ffun,1/3,[],[],[],[],-2,2,[],options)
```

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints were satisfied to within the default value of the constraint tolerance.

No active inequalities.

```
xfinal =
    -1.5056e-008
```

The true minimum occurs at $x = -2$. `fmincon` gives this report because the function $f(x)$ is so flat near $x = 0$.

Another common problem is that a solver finds a local minimum, but you might want a global minimum. For more information, see “Local vs. Global Optima” on page 4-27.

Lesson: check your results, even if the solver reports that it “found” a local minimum, or “solved” an equation.

This section gives techniques for verifying results.

1. Change the Initial Point

The initial point can have a large effect on the solution. If you obtain the same or worse solutions from various initial points, you become more confident in your solution.

For example, minimize $f(x) = x^3 + x^4$ starting from the point 1/4:

```
ffun = @(x)x^3 + x^4;  
options = optimoptions('fminunc','Algorithm','quasi-newton');  
[xfinal fval] = fminunc(ffun,1/4,options)
```

Local minimum found.

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```
x =  
-1.6764e-008
```

```
fval =  
-4.7111e-024
```

Change the initial point by a small amount, and the solver finds a better solution:

```
[xfinal fval] = fminunc(ffun,1/4+.001,options)
```

Local minimum found.

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```
xfinal =  
-0.7500
```

```
fval =
    -0.1055
```

$x = -0.75$ is the global solution; starting from other points cannot improve the solution.

For more information, see “Local vs. Global Optima” on page 4-27.

2. Check Nearby Points

To see if there are better values than a reported solution, evaluate your objective function and constraints at various nearby points.

For example, with the objective function `ffun` from “What Can Be Wrong If The Solver Succeeds?” on page 4-23, and the final point `xfinal = -1.5056e-008`, calculate `ffun(xfinal- δ)` for some δ :

```
delta = .1;
[ffun(xfinal),ffun(xfinal+delta),ffun(xfinal-delta)]

ans =
    -0.0000    0.0011   -0.0009
```

The objective function is lower at `ffun(xfinal- δ)`, so the solver reported an incorrect solution.

A less trivial example:

```
options = optimoptions(@fmincon,'Algorithm','active-set');
lb = [0,-1]; ub = [1,1];
ffun = @(x)(x(1)-(x(1)-x(2))^2);
[x fval exitflag] = fmincon(ffun,[1/2 1/3],[[],[]],[[],[]],...
                           lb,ub,[],options)
```

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints were satisfied to within the default value of the constraint tolerance.

```
Active inequalities (to within options.TolCon = 1e-006):
```

```
  lower      upper      ineqlin      ineqnonlin
      1
```

```
x =
```

```
  1.0e-007 *
      0      0.1614
```

```
fval =
```

```
-2.6059e-016
```

```
exitflag =
```

```
1
```

Evaluating `ffun` at nearby feasible points shows that the solution `x` is not a true minimum:

```
[ffun([0, .001]), ffun([0, -.001]), ...
  ffun([.001, -.001]), ffun([.001, .001])]
```

```
ans =
```

```
  1.0e-003 *
 -0.0010  -0.0010   0.9960   1.0000
```

The first two listed values are smaller than the computed minimum `fval`.

If you have a Global Optimization Toolbox license, you can use the `patternsearch` function to check nearby points.

3. Check your Objective and Constraint Functions

Double-check your objective function and constraint functions to ensure that they correspond to the problem you intend to solve. Suggestions:

- Check the evaluation of your objective function at a few points.
- Check that each inequality constraint has the correct sign.
- If you performed a maximization, remember to take the negative of the reported solution. (This advice assumes that you maximized a function

by minimizing the negative of the objective.) For example, to maximize $f(x) = x - x^2$, minimize $g(x) = -x + x^2$:

```
options = optimoptions('fminunc','Algorithm','quasi-newton');
[x fval] = fminunc(@(x)-x+x^2,0,options)
```

Local minimum found.

Optimization completed because the size of the gradient is less than the default value of the function tolerance.

```
x =
    0.5000
```

```
fval =
   -0.2500
```

The maximum of f is 0.25, the negative of $fval$.

- Check that an infeasible point does not cause an error in your functions; see “Iterations Can Violate Constraints” on page 2-34.


Local vs. Global Optima

- “Why Didn’t the Solver Find the Smallest Minimum?” on page 4-27
- “Searching for a Smaller Minimum” on page 4-28
- “Basins of Attraction” on page 4-29

Why Didn’t the Solver Find the Smallest Minimum?

In general, solvers return a local minimum. The result might be a global minimum, but there is no guarantee that it is. This section describes why solvers behave this way, and gives suggestions for ways to search for a global minimum, if needed.

- A *local* minimum of a function is a point where the function value is smaller than at nearby points, but possibly greater than at a distant point.
- A *global* minimum is a point where the function value is smaller than at all other feasible points.



Generally, Optimization Toolbox solvers find a local optimum. (This local optimum can be a global optimum.) They find the optimum in the *basin of attraction* of the starting point. For more information about basins of attraction, see “Basins of Attraction” on page 4-29.

There are some exceptions to this general rule.

- Linear programming and positive definite quadratic programming problems are convex, with convex feasible regions, so there is only one basin of attraction. Indeed, under certain choices of options, `linprog` ignores any user-supplied starting point, and `quadprog` does not require one, though supplying one can sometimes speed a minimization.
- Global Optimization Toolbox functions, such as `simulannealbnd`, attempt to search more than one basin of attraction.

Searching for a Smaller Minimum

If you need a global optimum, you must find an initial value for your solver in the basin of attraction of a global optimum.

Suggestions for ways to set initial values to search for a global optimum:

- Use a regular grid of initial points.
- Use random points drawn from a uniform distribution if your problem has all its coordinates bounded. Use points drawn from normal, exponential, or other random distributions if some components are unbounded. The less you know about the location of the global optimum, the more spread-out your random distribution should be. For example, normal distributions rarely sample more than three standard deviations away from their means, but a Cauchy distribution (density $1/(1+x^2)$) makes hugely disparate samples.

- Use identical initial points with added random perturbations on each coordinate, bounded, normal, exponential, or other.
- If you have a Global Optimization Toolbox license, use the `GlobalSearch` or `MultiStart` solvers. These solvers automatically generate random start points within bounds.

The more you know about possible initial points, the more focused and successful your search will be.


Basins of Attraction

If an objective function $f(x)$ is smooth, the vector $-\nabla f(x)$ points in the direction where $f(x)$ decreases most quickly. The equation of steepest descent, namely

$$\frac{d}{dt}x(t) = -\nabla f(x(t)),$$


yields a path $x(t)$ that goes to a local minimum as t gets large. Generally, initial values $x(0)$ that are near to each other give steepest descent paths that tend to the same minimum point. The *basin of attraction* for steepest descent is the set of initial values that lead to the same local minimum.

The following figure shows two one-dimensional minima. The figure shows different basins of attraction with different line styles, and shows directions of steepest descent with arrows. For this and subsequent figures, black dots represent local minima. Every steepest descent path, starting at a point $x(0)$, goes to the black dot in the basin containing $x(0)$.




One-dimensional basins

The following figure shows how steepest descent paths can be more complicated in more dimensions.



One basin of attraction, showing steepest descent paths from various starting points

The following figure shows even more complicated paths and basins of attraction.




Several basins of attraction

Constraints can break up one basin of attraction into several pieces. For example, consider minimizing y subject to:

- $y \geq |x|$
- $y \geq 5 - 4(x-2)^2$.

The figure shows the two basins of attraction with the final points.



The steepest descent paths are straight lines down to the constraint boundaries. From the constraint boundaries, the steepest descent paths travel down along the boundaries. The final point is either $(0,0)$ or $(11/4,11/4)$, depending on whether the initial x -value is above or below 2.

Optimizing a Simulation or Ordinary Differential Equation

In this section...

“What Is Optimizing a Simulation or ODE?” on page 4-33

“Potential Problems and Solutions” on page 4-33

“Bibliography” on page 4-39

What Is Optimizing a Simulation or ODE?

Sometimes your objective function or nonlinear constraint function values are available only by simulation or by numerical solution of an ordinary differential equation (ODE). Such optimization problems have several common characteristics and challenges, discussed in “Potential Problems and Solutions” on page 4-33.

To optimize a Simulink® model easily, try using Simulink Design Optimization™.

Potential Problems and Solutions

- “Problems in Finite Differences” on page 4-33
- “Suggestions for Finite Differences” on page 4-34
- “Problems in Stochastic Functions” on page 4-37
- “Suggestions for Stochastic Functions” on page 4-37
- “Common Calculation of Objective and Constraints” on page 4-37
- “Failure in Objective or Constraint Function Evaluation” on page 4-38
- “Suggestions for Evaluation Failures” on page 4-38

Problems in Finite Differences

Optimization Toolbox solvers use derivatives of objective and constraint functions internally. By default, they estimate these derivatives using finite difference approximations of the form

$$\frac{F(x + \delta) - F(x)}{\delta}$$

or

$$\frac{F(x + \delta) - F(x - \delta)}{2\delta}.$$

These finite difference approximations can be inaccurate because:

- A large value of δ allows more nonlinearity to affect the finite difference.
- A small value of δ leads to inaccuracy due to limited precision in numerics.

Specifically, for simulations and numerical solutions of ODEs:

- Simulations are often insensitive to small changes in parameters. This means that if you use too small a perturbation δ , the simulation can return a spurious estimated derivative of 0.
- Both simulations and numerical solutions of ODEs can have inaccuracies in their function evaluations. These inaccuracies can be amplified in finite difference approximations.
- Numerical solution of ODEs introduces noise at values much larger than machine precision. This noise can be amplified in finite difference approximations.
- If an ODE solver uses variable step sizes, then sometimes the number of ODE steps in the evaluation of $F(x + \delta)$ can differ from the number of steps in the evaluation of $F(x)$. This difference can lead to a spurious difference in the returned values, giving a misleading estimate of the derivative.

Suggestions for Finite Differences

- “Avoid Finite Differences by Using Direct Search” on page 4-35
- “Set Larger Finite Differences” on page 4-35
- “Use a Gradient Evaluation Function” on page 4-36
- “Use Tighter ODE Tolerances” on page 4-37

Avoid Finite Differences by Using Direct Search. If you have a Global Optimization Toolbox license, you can try using `patternsearch` as your solver. `patternsearch` does not attempt to estimate gradients, so does not suffer from the limitations in “Problems in Finite Differences” on page 4-33.

If you use `patternsearch` for expensive (time-consuming) function evaluations, use the `Cache` option:

```
options = psoptimset('Cache','on');
```

If you cannot use `patternsearch`, and have a relatively low-dimensional unconstrained minimization problem, try `fminsearch` instead. `fminsearch` does not use finite differences. However, `fminsearch` is not a fast or tunable solver.

Set Larger Finite Differences. You can sometimes avoid the problems in “Problems in Finite Differences” on page 4-33 by taking larger finite difference steps than the default.

- If you have MATLAB R2011b or later, set the `FinDiffRelStep` option to a value larger than the default `sqrt(eps)` or `eps^(1/3)`, such as:

- For R2011b–R2012b:

```
options = optimset('FinDiffRelStep',1e-3);
```

- For R2013a onwards:

```
options = optimoptions('solvername','FinDiffRelStep',1e-3);
```

If you have different scales in different components, set `FinDiffRelStep` to a vector proportional to the component scales.

- If you have MATLAB R2011a or earlier, set the `DiffMinChange` option to a larger value than the default `1e-8`, and possibly set the `DiffMaxChange` option also, such as:

```
options = optimset('DiffMinChange',1e-3,'DiffMaxChange',1);
```

Note It is difficult to know how to set these finite difference sizes.

You can also try setting central finite differences:

```
options = optimoptions('solvername','FinDiffType','central');
```

Use a Gradient Evaluation Function. To avoid the problems of finite difference estimation, you can give an approximate gradient function in your objective and nonlinear constraints. Remember to set the `GradObj` option to 'on' using `optimoptions`, and, if relevant, also set the `GradConstr` option to 'on'.

- For some ODEs, you can evaluate the gradient numerically at the same time as you solve the ODE. For example, suppose the differential equation for your objective function $z(t,x)$ is

$$\frac{d}{dt}z(t,x) = G(z,t,x),$$

where x is the vector of parameters over which you minimize. Suppose x is a scalar. Then the differential equation for its derivative y ,

$$y(t,x) = \frac{d}{dx}z(t,x)$$

is

$$\frac{d}{dt}y(t,x) = \frac{\partial G(z,t,x)}{\partial z}y(t,x) + \frac{\partial G(z,t,x)}{\partial x},$$

where $z(t,x)$ is the solution of the objective function ODE. You can solve for $y(t,x)$ in the same system of differential equations as $z(t,x)$. This solution gives you an approximated derivative without ever taking finite differences. For nonscalar x , solve one ODE per component.

For theoretical and computational aspects of this method, see Leis and Kramer [2]. For computational experience with this and finite-difference methods, see Figure 7 of Raue et al. [3].

- For some simulations, you can estimate a derivative within the simulation. For example, the likelihood ratio technique described in Reiman and Weiss [4] or the infinitesimal perturbation analysis technique analyzed in

Heidelberger, Cao, Zazanis, and Suri [1] estimate derivatives in the same simulation that estimates the objective or constraint functions.

Use Tighter ODE Tolerances. You can use `odeset` to set the `AbsTol` or `RelTol` ODE solver tolerances to values below their defaults. However, choosing too small a tolerance can lead to slow solutions, convergence failure, or other problems. Never choose a tolerance less than $1e-9$ for `RelTol`. The lower limit on each component of `AbsTol` depends on the scale of your problem, so there is no advice.

Problems in Stochastic Functions

If a simulation uses random numbers, then evaluating an objective or constraint function twice can return different results. This affects both function estimation and finite difference estimation. The value of a finite difference might be dominated by the variation due to randomness, rather than the variation due to different evaluation points x and $x + \delta$.

Suggestions for Stochastic Functions

If your simulation uses random numbers from a stream you control, reset the random stream before each evaluation of your objective or constraint functions. This practice can reduce the variability in results. For example, in an objective function:

```
function f = mysimulation(x)
rng('default') % or any other resetting method
...
end
```

For details, see “Generate Random Numbers That Are Repeatable”.

Common Calculation of Objective and Constraints

Frequently, a simulation evaluates both the objective function and constraints during the same simulation run. Or, both objective and nonlinear constraint functions use the same expensive computation. Solvers such as `fmincon` separately evaluate the objective function and nonlinear constraint functions. This can lead to a great loss of efficiency, because the solver calls the expensive computation twice. To circumvent this problem, use the technique in “Objective and Nonlinear Constraints in the Same Function” on page 2-48.

Failure in Objective or Constraint Function Evaluation

Your simulation or ODE can fail for some parameter values.

Suggestions for Evaluation Failures

Set Appropriate Bounds. While you might not know all limitations on the parameter space, try to set appropriate bounds on all parameters, both upper and lower. This can speed up your optimization, and can help the solver avoid problematic parameter values.

Use a Solver That Respects Bounds. As described in “Iterations Can Violate Constraints” on page 2-34, some algorithms can violate bound constraints at intermediate iterations. For optimizing simulations and ODEs, use algorithms that always obey bound constraints. See “Algorithms That Satisfy Bound Constraints” on page 2-34.

Return NaN. If your simulation or ODE solver does not successfully evaluate an objective or nonlinear constraint function at a point x , have your function return NaN. Most Optimization Toolbox and Global Optimization Toolbox solvers have the robustness to attempt a different iterative step if they encounter a NaN value. These robust solvers include:

- `fmincon` interior-point, `sqp`, and `trust-region-reflective` algorithms
- `fminunc`
- `lsqcurvefit`
- `lsqnonlin`
- `patternsearch`

Some people are tempted to return an arbitrary large objective function value at an unsuccessful, infeasible, or other poor point. However, this practice can confuse a solver, because the solver does not realize that the returned value is arbitrary. When you return NaN, the solver can attempt to evaluate at a different point.

Bibliography

- [1] Heidelberger, P., X.-R. Cao, M. A. Zazanis, and R. Suri. *Convergence properties of infinitesimal perturbation analysis estimates*. Management Science 34, No. 11, pp. 1281–1302, 1988.
- [2] Leis, J. R., and Kramer, M.A. *The Simultaneous Solution and Sensitivity Analysis of Systems Described by Ordinary Differential Equations*. ACM Trans. Mathematical Software, Vol. 14, No. 1, pp. 45–60, 1988.
- [3] Raue, A., et al. *Lessons Learned from Quantitative Dynamical Modeling in Systems Biology*. Available at <http://www.plosone.org/article/info:doi/10.1371/journal.pone.0074335>, 2013.
- [4] Reiman, M. I., and A. Weiss. *Sensitivity analysis via likelihood ratios*. Proc. 18th Winter Simulation Conference, ACM, New York, pp. 285–289, 1986.

Optimization App

Optimization App

In this section...
“Optimization App Basics” on page 5-2
“Specifying Certain Options” on page 5-10
“Importing and Exporting Your Work” on page 5-13

Optimization App Basics


- “How to Open the Optimization App” on page 5-2
- “Steps for Using the Optimization App” on page 5-4
- “Pausing and Stopping” on page 5-5
- “Viewing Results” on page 5-5
- “Final Point” on page 5-8
- “Starting a New Problem” on page 5-9

How to Open the Optimization App


To open the Optimization app, type

```
optimtool
```

in the Command Window. This opens the Optimization app, as shown in the following figure.



You can also start the Optimization app from the MATLAB **Apps** tab.



The reference page for the Optimization app provides variations for starting the `optimtool` function.

Steps for Using the Optimization App

This is a summary of the steps to set up your optimization problem and view results with the Optimization app.

The screenshot shows the Optimization Tool app interface with the following components and annotations:

- 1. Select solver and algorithm:** Points to the Solver dropdown menu (set to 'fmincon - Constrained nonlinear minimization') and the Algorithm dropdown menu (set to 'Trust region reflective').
- 2. Specify function to minimize:** Points to the Objective function dropdown menu.
- 3. Set problem parameters for selected solver:** Points to the Constraints section, including Linear inequalities (A, b), Linear equalities (Aeq, beq), Bounds (Lower, Upper), and Nonlinear constraint function.
- 4. Specify options:** Points to the Options panel on the right, which includes Stopping criteria (Max iterations, Max function evaluations, X tolerance, Function tolerance, Nonlinear constraint tolerance, SQP constraint tolerance, Unboundedness threshold) and Function value check options.
- 5. Run solver:** Points to the Start, Pause, and Stop buttons.
- 6. View solver status and results:** Points to the Current iteration field and the Clear Results button.
- 7. Import and export problem, options, and results:** Points to the File and Help menus at the top of the window.

Pausing and Stopping

While a solver is running, you can

- Click **Pause** to temporarily suspend the algorithm. To resume the algorithm using the current iteration at the time you paused, click **Resume**.
- Click **Stop** to stop the algorithm. The **Run solver and view results** window displays information for the current iteration at the moment you clicked **Stop**.

You can export your results after stopping the algorithm. For details, see “Exporting Your Work” on page 5-13.

Viewing Results

When a solver terminates, the **Run solver and view results** window displays the reason the algorithm terminated. To clear the **Run solver and view results** window between runs, click **Clear Results**.

Sorting the Displayed Results. Depending on the solver and problem, results can be in the form of a table. If the table has multiple rows, sort the table by clicking a column heading. Click the heading again to sort the results in reverse.

For example, suppose you use the Optimization app to solve the `bintprog` problem described in “Optimal Investments Via Binary Integer Programming” on page 6-151. The result appears as follows.

Run solver and view results

Start Pause Stop

Current iteration: 58 Clear Results

 Optimization running.
 Objective function value: -70.0
 Optimization terminated.

Final point:

Index	Value ▾
1	0
2	0
3	1
4	1
5	0
6	0
7	0
8	0
9	0
10	0
11	0
12	1
13	0
14	1
15	1
16	1
17	1
18	0
19	0

To sort the results by value, from lowest to highest, click **Value**.

Final point:	
Index	Value Δ
1	0
2	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
13	0
18	0
19	0
3	1
4	1
12	1
14	1
15	1
16	1
17	1


To sort the results in reverse order, highest to lowest, click **Value** again.

Final point:	
Index	Value ▾
3	1
4	1
12	1
14	1
15	1
16	1
17	1
1	0
2	0
5	0
6	0
7	0
8	0
9	0
10	0
11	0
13	0
18	0
19	0

To return to the original order, click **Index**. For an example of sorting a table returned by the Global Optimization Toolbox `gamultiobj` function, see “Multiobjective Optimization with Two Objectives”.

If you export results using **File > Export to Workspace**, the exported results do not depend on the sorted display.

Final Point

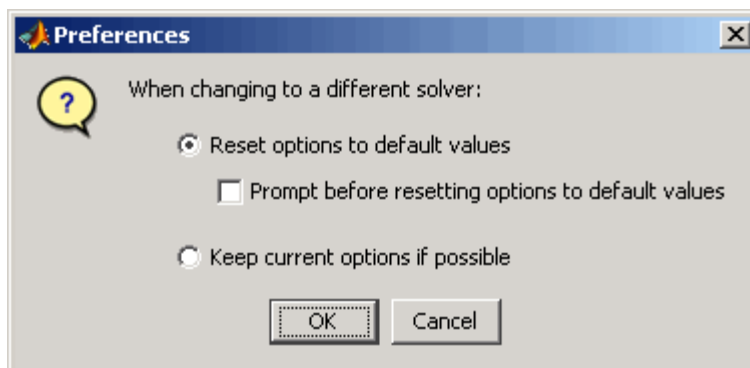
The **Final point** updates to show the coordinates of the final point when the algorithm terminated. If you don't see the final point, click the upward-pointing triangle on the  icon on the lower-left.

Starting a New Problem

Resetting Options and Clearing the Problem. Selecting **File > Reset Optimization Tool** resets the problem definition and options to the original default values. This action is equivalent to closing and restarting the app.

To clear only the problem definition, select **File > Clear Problem Fields**. With this action, fields in the **Problem Setup and Results** pane are reset to the defaults, with the exception of the selected solver and algorithm choice. Any options that you have modified from the default values in the **Options** pane are not reset with this action.

Setting Preferences for Changing Solvers. To modify how your options are handled in the Optimization app when you change solvers, select **File > Preferences**, which opens the Preferences dialog box shown below.



The default value, **Reset options to defaults**, discards any options you specified previously in the `optimtool`. Under this choice, you can select the option **Prompt before resetting options to defaults**.

Alternatively, you can select **Keep current options if possible** to preserve the values you have modified. Changed options that are not valid with the newly selected solver are kept but not used, while active options relevant to the new solver selected are used. This choice allows you to try different solvers with your problem without losing your options.

Specifying Certain Options

- “Plot Functions” on page 5-10
- “Output function” on page 5-11
- “Display to Command Window” on page 5-11


Plot Functions

You can select a plot function to easily plot various measures of progress while the algorithm executes. Each plot selected draws a separate axis in the figure window. If available for the solver selected, the **Stop** button in the **Run solver and view results** window to interrupt a running solver. You can select a predefined plot function from the Optimization app, or you can select **Custom function** to write your own. Plot functions not relevant to the solver selected are grayed out. The following lists the available plot functions:

- **Current point** — Select to show a bar plot of the point at the current iteration.
- **Function count** — Select to plot the number of function evaluations at each iteration.
- **Function value** — Select to plot the function value at each iteration.
- **Norm of residuals** — Select to show a bar plot of the current norm of residuals at the current iteration.
- **Max constraint** — Select to plot the maximum constraint violation value at each iteration.
- **Current step** — Select to plot the algorithm step size at each iteration.
- **First order optimality** — Select to plot the violation of the optimality conditions for the solver at each iteration.
- **Custom function** — Enter your own plot function as a function handle. To provide more than one plot function use a cell array, for example, by typing:

```
{@plotfcn,@plotfcn2}
```

Write custom plot functions with the same syntax as output functions. For information, see “Output Function” on page 9-21.



The graphic above shows the plot functions available for the default `fmincon` solver.

Output function

Output function is a function or collection of functions the algorithm calls at each iteration. Through an output function you can observe optimization quantities such as function values, gradient values, and current iteration. Specify no output function, a single output function using a function handle, or multiple output functions. To provide more than one output function use a cell array of function handles in the **Custom function** field, for example by typing:

```
{@outputfcn,@outputfcn2}
```

For more information on writing an output function, see “Output Function” on page 9-21.



Display to Command Window

Select **Level of display** to specify the amount of information displayed when you run the algorithm. Choose from the following; depending on the solver, only some may be available:

- `off` (default) — Display no output.
- `final` — Display the reason for stopping at the end of the run.

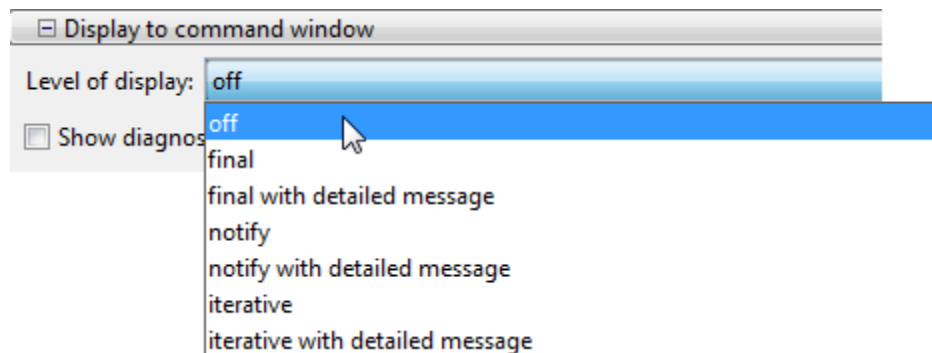
- **final with detailed message** — Display the detailed reason for stopping at the end of the run.
- **notify** — Display output only if the function does not converge.
- **notify with detailed message** — Display a detailed output only if the function does not converge.
- **iterative** — Display information at each iteration of the algorithm and the reason for stopping at the end of the run.
- **iterative with detailed message** — Display information at each iteration of the algorithm and the detailed reason for stopping at the end of the run.

See “Enhanced Exit Messages” on page 3-5 for information on detailed messages.

Set **Node interval**, with the `bintprog` solver selected, to specify the interval of explored nodes you want to display output for. Note that integer feasible solution nodes are always shown.

Selecting **Show diagnostics** lists problem information and options that have changed from the defaults.

The graphic below shows the display options for the `fmincon` solver. Some other solvers have fewer options.



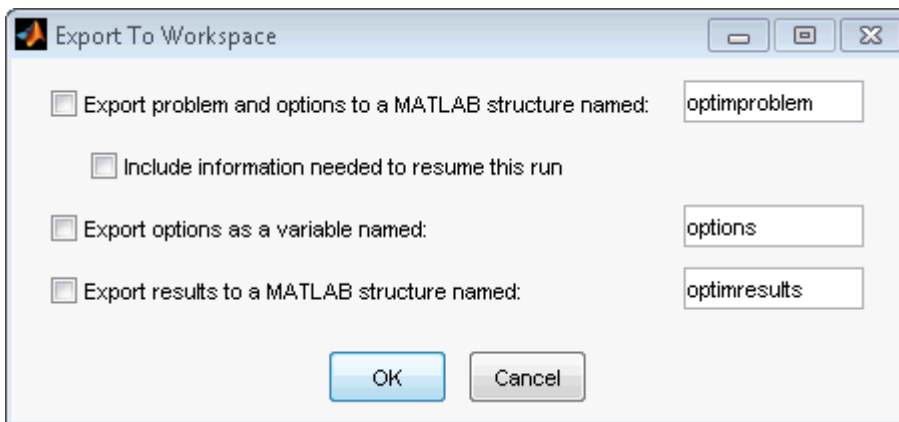
Importing and Exporting Your Work

- “Exporting Your Work” on page 5-13
- “Importing Your Work” on page 5-15
- “Generating a File” on page 5-15

Exporting Your Work

The **Export to Workspace** dialog box enables you to send your problem information to the MATLAB workspace as a structure or object that you may then manipulate in the Command Window.

To access the **Export to Workspace** dialog box shown below, select **File > Export to Workspace**.



You can specify results that contain:

- The problem and options information
- The problem and options information, and the state of the solver when stopped (this means the latest point for most solvers, the current population for Genetic Algorithms solvers, and the best point found for the Simulated Annealing solver)

- The states of random number generators `rand` and `randn` at the start of the previous run, by checking the **Use random states from previous run** box for applicable solvers
- The options information only
- The results of running your problem in the Optimization app

Exported results contain all optional information. For example, an exported results structure for `lsqcurvefit` contains the data `x`, `resnorm`, `residual`, `exitflag`, `output`, `lambda`, and `jacobian`.

After you have exported information from the Optimization app to the MATLAB workspace, you can see your data in the MATLAB Workspace browser or by typing the name of the structure at the Command Window. To see the value of a field in a structure or object, double-click the name in the Workspace window. Alternatively, see the values by entering `exportname.fieldname` at the command line. For example, to see the message in an output structure, enter `output.message`. If a structure contains structures or objects, you can double-click again in the workspace browser, or enter `exportname.name2.fieldname` at the command line. For example, to see the level of iterative display contained in the options of an exported problem structure, enter `optimproblem.options.Display`.

You can run a solver on an exported problem at the command line by typing

```
solver(problem)
```

For example, if you have exported a `fmincon` problem named `optimproblem`, you can type

```
fmincon(optimproblem)
```

This runs `fmincon` on the problem with the saved options in `optimproblem`. You can exercise more control over outputs by typing, for example,

```
[x,fval,exitflag] = fmincon(optimproblem)
```

or use any other supported syntax.

Importing Your Work

Whether you save options from Optimization Toolbox functions at the Command Window, or whether you export options, or the problem and options, from the Optimization app, you can resume work on your problem using the Optimization app.

There are three ways to import your options, or problem and options, to the Optimization app:

- Call the `optimtool` function from the Command Window specifying your options, or problem and options, as the input, for example,

```
optimtool(options)
```

- Select **File > Import Options** in the Optimization app.
- Select **File > Import Problem** in the Optimization app.

The methods described above require that the options, or problem and options, be present in the MATLAB workspace.

If you import a problem that was generated with the **Include information needed to resume this run** box checked, the initial point is the latest point generated in the previous run. (For Genetic Algorithm solvers, the initial population is the latest population generated in the previous run. For the Simulated Annealing solver, the initial point is the best point generated in the previous run.) If you import a problem that was generated with this box unchecked, the initial point (or population) is the initial point (or population) of the previous run.

Generating a File

You may want to generate a file to continue with your optimization problem in the Command Window at another time. You can run the file without modification to recreate the results that you created with the Optimization app. You can also edit and modify the file and run it from the Command Window.

To export data from the Optimization app to a file, select **File > Generate Code**.

The generated file captures the following:

- The problem definition, including the solver, information on the function to be minimized, algorithm specification, constraints, and start point
- The options with the currently selected option value

Running the file at the Command Window reproduces your problem results.

Although you cannot export your problem results to a generated file, you can save them in a MAT-file that you can use with your generated file, by exporting the results using the Export to Workspace dialog box, then saving the data to a MAT-file from the Command Window.

Optimization Algorithms and Examples

- “Optimization Theory Overview” on page 6-4
- “Unconstrained Nonlinear Optimization Algorithms” on page 6-5
- “fminunc Unconstrained Minimization” on page 6-17
- “Minimization with Gradient and Hessian” on page 6-20
- “Minimization with Gradient and Hessian Sparsity Pattern” on page 6-22
- “Constrained Nonlinear Optimization Algorithms” on page 6-26
- “Nonlinear Inequality Constraints” on page 6-52
- “Nonlinear Constraints with Gradients” on page 6-54
- “fmincon Interior-Point Algorithm with Analytic Hessian” on page 6-57
- “Nonlinear Equality and Inequality Constraints” on page 6-65
- “Optimization App with the fmincon Solver” on page 6-67
- “Minimization with Bound Constraints and Banded Preconditioner” on page 6-72
- “Minimization with Linear Equality Constraints” on page 6-77
- “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80
- “Symbolic Math Toolbox Calculates Gradients and Hessians” on page 6-85
- “One-Dimensional Semi-Infinite Constraints” on page 6-100
- “Two-Dimensional Semi-Infinite Constraint” on page 6-103

- “Linear Programming Algorithms” on page 6-106
- “Typical Linear Programming Problem” on page 6-119
- “Quadratic Programming Algorithms” on page 6-121
- “Quadratic Minimization with Bound Constraints” on page 6-135
- “Quadratic Minimization with Dense, Structured Hessian” on page 6-138
- “Large Sparse Quadratic Program with Interior Point Algorithm” on page 6-144
- “Binary Integer Programming Algorithms” on page 6-148
- “Optimal Investments Via Binary Integer Programming” on page 6-151
- “Mixed-Integer Linear Programming Algorithms” on page 6-157
- “Tuning Integer Linear Programming” on page 6-166
- “Mixed-Integer Linear Programming Basics” on page 6-169
- “Factory, Warehouse, Sales Allocation Model” on page 6-173
- “Travelling Salesman Problem” on page 6-185
- “Solve Sudoku Puzzles Via Integer Programming” on page 6-192
- “Least-Squares (Model Fitting) Algorithms” on page 6-200
- “lsqnonlin with a Simulink Model” on page 6-209
- “Nonlinear Least Squares With and Without Jacobian” on page 6-215
- “Linear Least Squares with Bound Constraints” on page 6-219
- “Optimization App with the lsqin Solver” on page 6-221
- “Jacobian Multiply Function with Linear Least Squares” on page 6-224
- “Nonlinear Curve Fitting with lsqcurvefit” on page 6-229
- “Fit a Model to Complex-Valued Data” on page 6-231
- “Multiobjective Optimization Algorithms” on page 6-236
- “Using fminimax with a Simulink Model” on page 6-242
- “Signal Processing Using fgoalattain” on page 6-246
- “Equation Solving Algorithms” on page 6-249

-
- “Nonlinear Equations with Analytic Jacobian” on page 6-257
 - “Nonlinear Equations with Finite-Difference Jacobian” on page 6-261
 - “Nonlinear Equations with Jacobian” on page 6-263
 - “Nonlinear Equations with Jacobian Sparsity Pattern” on page 6-266
 - “Nonlinear Systems with Constraints” on page 6-269
 - “Bibliography” on page 6-274

Optimization Theory Overview

Optimization techniques are used to find a set of design parameters, $x = \{x_1, x_2, \dots, x_n\}$, that can in some way be defined as optimal. In a simple case this might be the minimization or maximization of some system characteristic that is dependent on x . In a more advanced formulation the objective function, $f(x)$, to be minimized or maximized, might be subject to constraints in the form of equality constraints, $G_i(x) = 0$ ($i = 1, \dots, m_e$); inequality constraints, $G_i(x) \leq 0$ ($i = m_e + 1, \dots, m$); and/or parameter bounds, x_l, x_u .

A General Problem (GP) description is stated as

$$\min_x f(x), \tag{6-1}$$

subject to

$$\begin{aligned} G_i(x) &= 0 & i = 1, \dots, m_e, \\ G_i(x) &\leq 0 & i = m_e + 1, \dots, m, \end{aligned}$$

where x is the vector of length n design parameters, $f(x)$ is the objective function, which returns a scalar value, and the vector function $G(x)$ returns a vector of length m containing the values of the equality and inequality constraints evaluated at x .

An efficient and accurate solution to this problem depends not only on the size of the problem in terms of the number of constraints and design variables but also on characteristics of the objective function and constraints. When both the objective function and the constraints are linear functions of the design variable, the problem is known as a Linear Programming (LP) problem. Quadratic Programming (QP) concerns the minimization or maximization of a quadratic objective function that is linearly constrained. For both the LP and QP problems, reliable solution procedures are readily available. More difficult to solve is the Nonlinear Programming (NP) problem in which the objective function and constraints can be nonlinear functions of the design variables. A solution of the NP problem generally requires an iterative procedure to establish a direction of search at each major iteration. This is usually achieved by the solution of an LP, a QP, or an unconstrained subproblem.

Unconstrained Nonlinear Optimization Algorithms

In this section...

“Unconstrained Optimization Definition” on page 6-5

“fminunc trust-region Algorithm” on page 6-5

“fminunc quasi-newton Algorithm” on page 6-8

“fminsearch Algorithm” on page 6-14

Unconstrained Optimization Definition

Unconstrained minimization is the problem of finding a vector x that is a local minimum to a scalar function $f(x)$:

$$\min_x f(x)$$

The term *unconstrained* means that no restriction is placed on the range of x .

fminunc trust-region Algorithm

Trust-Region Methods for Nonlinear Minimization

Many of the methods used in Optimization Toolbox solvers are based on *trust regions*, a simple yet powerful concept in optimization.

To understand the trust-region approach to optimization, consider the unconstrained minimization problem, minimize $f(x)$, where the function takes vector arguments and returns scalars. Suppose you are at a point x in n -space and you want to improve, i.e., move to a point with a lower function value. The basic idea is to approximate f with a simpler function q , which reasonably reflects the behavior of function f in a neighborhood N around the point x . This neighborhood is the trust region. A trial step s is computed by minimizing (or approximately minimizing) over N . This is the trust-region subproblem,

$$\min_s \{q(s), s \in N\}. \quad (6-2)$$

The current point is updated to be $x + s$ if $f(x + s) < f(x)$; otherwise, the current point remains unchanged and N , the region of trust, is shrunk and the trial step computation is repeated.

The key questions in defining a specific trust-region approach to minimizing $f(x)$ are how to choose and compute the approximation q (defined at the current point x), how to choose and modify the trust region N , and how accurately to solve the trust-region subproblem. This section focuses on the unconstrained problem. Later sections discuss additional complications due to the presence of constraints on the variables.

In the standard trust-region method ([48]), the quadratic approximation q is defined by the first two terms of the Taylor approximation to F at x ; the neighborhood N is usually spherical or ellipsoidal in shape. Mathematically the trust-region subproblem is typically stated

$$\min \left\{ \frac{1}{2} s^T H s + s^T g \text{ such that } \|D s\| \leq \Delta \right\}, \quad (6-3)$$

where g is the gradient of f at the current point x , H is the Hessian matrix (the symmetric matrix of second derivatives), D is a diagonal scaling matrix, Δ is a positive scalar, and $\| \cdot \|$ is the 2-norm. Good algorithms exist for solving Equation 6-3 (see [48]); such algorithms typically involve the computation of a full eigensystem and a Newton process applied to the secular equation

$$\frac{1}{\Delta} - \frac{1}{\|s\|} = 0.$$

Such algorithms provide an accurate solution to Equation 6-3. However, they require time proportional to several factorizations of H . Therefore, for large-scale problems a different approach is needed. Several approximation and heuristic strategies, based on Equation 6-3, have been proposed in the literature ([42] and [50]). The approximation approach followed in Optimization Toolbox solvers is to restrict the trust-region subproblem to a two-dimensional subspace S ([39] and [42]). Once the subspace S has been computed, the work to solve Equation 6-3 is trivial even if full

eigenvalue/eigenvector information is needed (since in the subspace, the problem is only two-dimensional). The dominant work has now shifted to the determination of the subspace.

The two-dimensional subspace S is determined with the aid of a preconditioned conjugate gradient process described below. The solver defines S as the linear space spanned by s_1 and s_2 , where s_1 is in the direction of the gradient g , and s_2 is either an approximate Newton direction, i.e., a solution to

$$H \cdot s_2 = -g, \quad (6-4)$$

or a direction of negative curvature,

$$s_2^T \cdot H \cdot s_2 < 0. \quad (6-5)$$

The philosophy behind this choice of S is to force global convergence (via the steepest descent direction or negative curvature direction) and achieve fast local convergence (via the Newton step, when it exists).

A sketch of unconstrained minimization using trust-region ideas is now easy to give:

- 1** Formulate the two-dimensional trust-region subproblem.
- 2** Solve Equation 6-3 to determine the trial step s .
- 3** If $f(x + s) < f(x)$, then $x = x + s$.
- 4** Adjust Δ .

These four steps are repeated until convergence. The trust-region dimension Δ is adjusted according to standard rules. In particular, it is decreased if the trial step is not accepted, i.e., $f(x + s) \geq f(x)$. See [46] and [49] for a discussion of this aspect.

Optimization Toolbox solvers treat a few important special cases of f with specialized functions: nonlinear least-squares, quadratic functions, and linear least-squares. However, the underlying algorithmic ideas are the same as for the general case. These special cases are discussed in later sections.

Preconditioned Conjugate Gradient Method

A popular way to solve large symmetric positive definite systems of linear equations $Hp = -g$ is the method of Preconditioned Conjugate Gradients (PCG). This iterative approach requires the ability to calculate matrix-vector products of the form Hv where v is an arbitrary vector. The symmetric positive definite matrix M is a *preconditioner* for H . That is, $M = C^2$, where $C^{-1}HC^{-1}$ is a well-conditioned matrix or a matrix with clustered eigenvalues.

In a minimization context, you can assume that the Hessian matrix H is symmetric. However, H is guaranteed to be positive definite only in the neighborhood of a strong minimizer. Algorithm PCG exits when a direction of negative (or zero) curvature is encountered, i.e., $d^T Hd \leq 0$. The PCG output direction, p , is either a direction of negative curvature or an approximate (*tol* controls how approximate) solution to the Newton system $Hp = -g$. In either case p is used to help define the two-dimensional subspace used in the trust-region approach discussed in “Trust-Region Methods for Nonlinear Minimization” on page 6-5.

fminunc quasi-newton Algorithm

Basics of Unconstrained Optimization

Although a wide spectrum of methods exists for unconstrained optimization, methods can be broadly categorized in terms of the derivative information that is, or is not, used. Search methods that use only function evaluations (e.g., the simplex search of Nelder and Mead [30]) are most suitable for problems that are not smooth or have a number of discontinuities. Gradient methods are generally more efficient when the function to be minimized is continuous in its first derivative. Higher order methods, such as Newton’s method, are only really suitable when the second-order information is readily and easily calculated, because calculation of second-order information, using numerical differentiation, is computationally expensive.

Gradient methods use information about the slope of the function to dictate a direction of search where the minimum is thought to lie. The simplest of these is the method of steepest descent in which a search is performed in a direction, $-\nabla f(x)$, where $\nabla f(x)$ is the gradient of the objective function. This method is very inefficient when the function to be minimized has long narrow valleys as, for example, is the case for Rosenbrock’s function

$$f(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2. \quad (6-6)$$

The minimum of this function is at $x = [1, 1]$, where $f(x) = 0$. A contour map of this function is shown in the figure below, along with the solution path to the minimum for a steepest descent implementation starting at the point $[-1.9, 2]$. The optimization was terminated after 1000 iterations, still a considerable distance from the minimum. The black areas are where the method is continually zigzagging from one side of the valley to another. Note that toward the center of the plot, a number of larger steps are taken when a point lands exactly at the center of the valley.




Figure 6-1: Steepest Descent Method on Rosenbrock's Function (Equation 6-6)

This function, also known as the banana function, is notorious in unconstrained examples because of the way the curvature bends around the origin. Rosenbrock's function is used throughout this section to illustrate the use of a variety of optimization techniques. The contours have been plotted in exponential increments because of the steepness of the slope surrounding the U-shaped valley.

For an animated version of this figure, enter `bandem` at the MATLAB command line.

Quasi-Newton Methods

Of the methods that use gradient information, the most favored are the quasi-Newton methods. These methods build up curvature information at each iteration to formulate a quadratic model problem of the form

$$\min_x \frac{1}{2} x^T H x + c^T x + b, \quad (6-7)$$

where the Hessian matrix, H , is a positive definite symmetric matrix, c is a constant vector, and b is a constant. The optimal solution for this problem occurs when the partial derivatives of x go to zero, i.e.,

$$\nabla f(x^*) = H x^* + c = 0. \quad (6-8)$$

The optimal solution point, x^* , can be written as

$$x^* = -H^{-1}c. \quad (6-9)$$

Newton-type methods (as opposed to quasi-Newton methods) calculate H directly and proceed in a direction of descent to locate the minimum after a number of iterations. Calculating H numerically involves a large amount of computation. Quasi-Newton methods avoid this by using the observed behavior of $f(x)$ and $\nabla f(x)$ to build up curvature information to make an approximation to H using an appropriate updating technique.

A large number of Hessian updating methods have been developed. However, the formula of Broyden [3], Fletcher [12], Goldfarb [20], and Shanno [37] (BFGS) is thought to be the most effective for use in a general purpose method.

The formula given by BFGS is

$$H_{k+1} = H_k + \frac{q_k q_k^T}{q_k^T s_k} - \frac{H_k s_k s_k^T H_k^T}{s_k^T H_k s_k}, \quad (6-10)$$

where

$$\begin{aligned}s_k &= x_{k+1} - x_k, \\ q_k &= \nabla f(x_{k+1}) - \nabla f(x_k).\end{aligned}$$

As a starting point, H_0 can be set to any symmetric positive definite matrix, for example, the identity matrix I . To avoid the inversion of the Hessian H , you can derive an updating method that avoids the direct inversion of H by using a formula that makes an approximation of the inverse Hessian H^{-1} at each update. A well-known procedure is the DFP formula of Davidon [7], Fletcher, and Powell [14]. This uses the same formula as the BFGS method (Equation 6-10) except that q_k is substituted for s_k .

The gradient information is either supplied through analytically calculated gradients, or derived by partial derivatives using a numerical differentiation method via finite differences. This involves perturbing each of the design variables, x , in turn and calculating the rate of change in the objective function.

At each major iteration, k , a line search is performed in the direction

$$d = -H_k^{-1} \cdot \nabla f(x_k). \quad (6-11)$$

The quasi-Newton method is illustrated by the solution path on Rosenbrock's function in Figure 6-2, BFGS Method on Rosenbrock's Function. The method is able to follow the shape of the valley and converges to the minimum after 140 function evaluations using only finite difference gradients.




Figure 6-2: BFGS Method on Rosenbrock's Function

For an animated version of this figure, enter `bandem` at the MATLAB command line.

Line Search

Line search is a search method that is used as part of a larger optimization algorithm. At each step of the main algorithm, the line-search method searches along the line containing the current point, x_k , parallel to the *search direction*, which is a vector determined by the main algorithm. That is, the method finds the next iterate x_{k+1} of the form

$$x_{k+1} = x_k + \alpha * d_k, \quad (6-12)$$

where x_k denotes the current iterate, d_k is the search direction, and α^* is a scalar step length parameter.

The line search method attempts to decrease the objective function along the line $x_k + \alpha^* d_k$ by repeatedly minimizing polynomial interpolation models of the objective function. The line search procedure has two main steps:

- The *bracketing* phase determines the range of points on the line $x_{k+1} = x_k + \alpha * d_k$ to be searched. The *bracket* corresponds to an interval specifying the range of values of α .

- The *sectioning* step divides the bracket into subintervals, on which the minimum of the objective function is approximated by polynomial interpolation.

The resulting step length α satisfies the Wolfe conditions:

$$f(x_k + \alpha d_k) \leq f(x_k) + c_1 \alpha \nabla f_k^T d_k, \quad (6-13)$$

$$\nabla f(x_k + \alpha d_k)^T d_k \geq c_2 \nabla f_k^T d_k, \quad (6-14)$$

where c_1 and c_2 are constants with $0 < c_1 < c_2 < 1$.

The first condition (Equation 6-13) requires that α_k sufficiently decreases the objective function. The second condition (Equation 6-14) ensures that the step length is not too small. Points that satisfy both conditions (Equation 6-13 and Equation 6-14) are called *acceptable points*.

The line search method is an implementation of the algorithm described in Section 2-6 of [13]. See also [31] for more information about line search.

Hessian Update

Many of the optimization functions determine the direction of search by updating the Hessian matrix at each iteration, using the BFGS method (Equation 6-10). The function `fminunc` also provides an option to use the DFP method given in “Quasi-Newton Methods” on page 6-10 (set `HessUpdate` to `'dfp'` in `options` to select the DFP method). The Hessian, H , is always maintained to be positive definite so that the direction of search, d , is always in a descent direction. This means that for some arbitrarily small step α in the direction d , the objective function decreases in magnitude. You achieve positive definiteness of H by ensuring that H is initialized to be positive

definite and thereafter $q_k^T s_k$ (from Equation 6-15) is always positive. The

term $q_k^T s_k$ is a product of the line search step length parameter α_k and a combination of the search direction d with past and present gradient evaluations,

$$q_k^T s_k = \alpha_k \left(\nabla f(x_{k+1})^T d - \nabla f(x_k)^T d \right). \quad (6-15)$$

You always achieve the condition that $q_k^T s_k$ is positive by performing a sufficiently accurate line search. This is because the search direction, d , is a descent direction, so that a_k and the negative gradient $-\nabla f(x_k)^T d$ are always positive. Thus, the possible negative term $-\nabla f(x_{k+1})^T d$ can be made as small in magnitude as required by increasing the accuracy of the line search.

fminsearch Algorithm

`fminsearch` uses the Nelder-Mead simplex algorithm as described in Lagarias et al. [57]. This algorithm uses a simplex of $n + 1$ points for n -dimensional vectors x . The algorithm first makes a simplex around the initial guess x_0 by adding 5% of each component $x_0(i)$ to x_0 , and using these n vectors as elements of the simplex in addition to x_0 . (It uses 0.00025 as component i if $x_0(i) = 0$.) Then, the algorithm modifies the simplex repeatedly according to the following procedure.

Note The keywords for the `fminsearch` iterative display appear in bold after the description of the step.

- 1** Let $x(i)$ denote the list of points in the current simplex, $i = 1, \dots, n+1$.
- 2** Order the points in the simplex from lowest function value $f(x(1))$ to highest $f(x(n+1))$. At each step in the iteration, the algorithm discards the current worst point $x(n+1)$, and accepts another point into the simplex. [Or, in the case of step 7 below, it changes all n points with values above $f(x(1))$].
- 3** Generate the *reflected* point

$$r = 2m - x(n+1),$$

where

$$m = \sum x(i)/n, \quad i = 1 \dots n,$$

and calculate $f(r)$.

4 If $f(x(1)) \leq f(r) < f(x(n))$, accept r and terminate this iteration. **Reflect**

5 If $f(r) < f(x(1))$, calculate the expansion point s

$$s = m + 2(m - x(n+1)),$$

and calculate $f(s)$.

a If $f(s) < f(r)$, accept s and terminate the iteration. **Expand**

b Otherwise, accept r and terminate the iteration. **Reflect**

6 If $f(r) \geq f(x(n))$, perform a *contraction* between m and the better of $x(n+1)$ and r :

a If $f(r) < f(x(n+1))$ (i.e., r is better than $x(n+1)$), calculate

$$c = m + (r - m)/2$$

and calculate $f(c)$. If $f(c) < f(r)$, accept c and terminate the iteration.

Contract outside Otherwise, continue with Step 7 (Shrink).

b If $f(r) \geq f(x(n+1))$, calculate

$$cc = m + (x(n+1) - m)/2$$


and calculate $f(cc)$. If $f(cc) < f(x(n+1))$, accept cc and terminate the iteration. **Contract inside** Otherwise, continue with Step 7 (Shrink).

7 Calculate the n points

$$v(i) = x(1) + (x(i) - x(1))/2$$

and calculate $f(v(i))$, $i = 2, \dots, n+1$. The simplex at the next iteration is $x(1)$, $v(2), \dots, v(n+1)$. **Shrink**

The following figure shows the points that `fminsearch` might calculate in the procedure, along with each possible new simplex. The original simplex has a bold outline. The iterations proceed until they meet a stopping criterion.



fminunc Unconstrained Minimization

Consider the problem of finding a set of values $[x_1, x_2]$ that solves

$$\min_x f(x) = e^{x_1} (4x_1^2 + 2x_2^2 + 4x_1x_2 + 2x_2 + 1). \quad (6-16)$$

To solve this two-dimensional problem, write a file that returns the function value. Then, invoke the unconstrained minimization routine `fminunc`.

Step 1: Write a file `objfun.m`.

This code ships with the toolbox. To view, enter type `objfun`:

```
function f = objfun(x)
f = exp(x(1)) * (4*x(1)^2 + 2*x(2)^2 + 4*x(1)*x(2) + 2*x(2) + 1);
```

Step 2: Invoke one of the unconstrained optimization routines.

```
x0 = [-1,1]; % Starting guess
options = optimoptions(@fminunc,'Algorithm','quasi-newton');
[x,fval,exitflag,output] = fminunc(@objfun,x0,options);
```

This produces the following output:

```
Local minimum found.
```

```
Optimization completed because the size of the gradient is less
than the default value of the function tolerance.
```

View the results:

```
x,fval,exitflag,output
```

```
x =
    0.5000   -1.0000
```

```
fval =
    3.6609e-15
```

```
exitflag =
    1

output =
    iterations: 8
    funcCount: 66
    stepsize: 1
firstorderopt: 1.2284e-007
    algorithm: 'medium-scale: Quasi-Newton line search'
    message: 'Local minimum found.'
```

Optimization completed because the size of the gradie...'

The `exitflag` tells whether the algorithm converged. `exitflag = 1` means a local minimum was found. The meanings of `exitflags` are given in function reference pages.

The output structure gives more details about the optimization. For `fminunc`, it includes the number of iterations in `iterations`, the number of function evaluations in `funcCount`, the final step-size in `stepsize`, a measure of first-order optimality (which in this unconstrained case is the infinity norm of the gradient at the solution) in `firstorderopt`, the type of algorithm used in `algorithm`, and the exit message (the reason the algorithm stopped).

Pass the variable `options` to `fminunc` to change characteristics of the optimization algorithm, as in

```
x = fminunc(@objfun,x0,options);
```

`options` contains values for termination tolerances and algorithm choices. Create `options` using the `optimoptions` function:

```
options = optimoptions(@fminunc,'Algorithm','quasi-newton');
```

You can also create `options` by exporting from the Optimization app.

In this example, we have used the `quasi-newton` algorithm. Other options include controlling the amount of command line display during the optimization iteration, the tolerances for the termination criteria, whether a user-supplied gradient or Jacobian is to be used, and the maximum number of iterations or function evaluations. See `optimoptions`, the individual

optimization functions, and “Optimization Options Reference” on page 9-7 for more options and information.

Minimization with Gradient and Hessian

This example involves solving a nonlinear minimization problem with a tridiagonal Hessian matrix $H(x)$ first computed explicitly, and then by providing the Hessian's sparsity structure for the finite-differencing routine.

The problem is to find x to minimize

$$f(x) = \sum_{i=1}^{n-1} \left((x_i^2)^{(x_{i+1}^2+1)} + (x_{i+1}^2)^{(x_i^2+1)} \right) \quad (6-17)$$

where $n = 1000$.

Step 1: Write a file `brownfgh.m` that computes the objective function, the gradient of the objective, and the sparse tridiagonal Hessian matrix.

The file is lengthy so is not included here. View the code with the command

```
type brownfgh
```

Because `brownfgh` computes the gradient and Hessian values as well as the objective function, you need to use `optimoptions` to indicate that this information is available in `brownfgh`, using the `GradObj` and `Hessian` options.

Step 2: Call a nonlinear minimization routine with a starting point `xstart`.

```
n = 1000;
xstart = -ones(n,1);
xstart(2:2:n,1) = 1;
options = optimoptions(@fminunc,'GradObj','on','Hessian','on');
[x,fval,exitflag,output] = fminunc(@brownfgh,xstart,options);
```

This 1000 variable problem is solved in about 7 iterations and 7 conjugate gradient iterations with a positive `exitflag` indicating convergence. The final function value and measure of optimality at the solution x are both close

to zero. For `fminunc`, the first order optimality is the infinity norm of the gradient of the function, which is zero at a local minimum:

```
fval,exitflag,output
```

```
fval =  
    2.8709e-17
```

```
exitflag =  
         1
```

```
output =  
    iterations: 7  
    funcCount: 8  
    cgiterations: 7  
    firstorderopt: 4.7948e-10  
    algorithm: 'large-scale: trust-region Newton'  
    message: 'Local minimum found.'
```

```
Optimization completed because the size of the grad...'  
    constrviolation: []
```

Minimization with Gradient and Hessian Sparsity Pattern

Next, solve the same problem but the Hessian matrix is now approximated by sparse finite differences instead of explicit computation. To use the trust-region method in `fminunc`, you *must* compute the gradient in `fun`; it is *not* optional as in the quasi-newton method.

The `brownfg` file computes the objective function and gradient.

Step 1: Write a file `brownfg.m` that computes the objective function and the gradient of the objective.

This function file ships with your software.

```
function [f,g] = brownfg(x)
% BROWNFG Nonlinear minimization test problem
%
% Evaluate the function
n=length(x); y=zeros(n,1);
i=1:(n-1);
y(i)=(x(i).^2).^(x(i+1).^2+1) + ...
      (x(i+1).^2).^(x(i).^2+1);
f=sum(y);
% Evaluate the gradient if nargout > 1
if nargout > 1
    i=1:(n-1); g = zeros(n,1);
    g(i) = 2*(x(i+1).^2+1).*x(i).* ...
           ((x(i).^2).^(x(i+1).^2))+ ...
           2*x(i).*((x(i+1).^2).^(x(i).^2+1)).* ...
           log(x(i+1).^2);
    g(i+1) = g(i+1) + ...
            2*x(i+1).*((x(i).^2).^(x(i+1).^2+1)).* ...
            log(x(i).^2) + ...
            2*(x(i).^2+1).*x(i+1).* ...
            ((x(i+1).^2).^(x(i).^2));
end
```


To allow efficient computation of the sparse finite-difference approximation of the Hessian matrix $H(x)$, the sparsity structure of H must be predetermined. In this case assume this structure, `Hstr`, a sparse matrix, is available in file

`brownhstr.mat`. Using the `spy` command you can see that `Hstr` is indeed sparse (only 2998 nonzeros). Use `optimoptions` to set the `HessPattern` option to `Hstr`. When a problem as large as this has obvious sparsity structure, not setting the `HessPattern` option requires a huge amount of unnecessary memory and computation because `fminunc` attempts to use finite differencing on a full Hessian matrix of one million nonzero entries.

You must also set the `GradObj` option to `'on'` using `optimoptions`, since the gradient is computed in `brownfg.m`. Then execute `fminunc` as shown in Step 2.

Step 2: Call a nonlinear minimization routine with a starting point `xstart`.

```
fun = @brownfg;
load brownhstr           % Get Hstr, structure of the Hessian
spy(Hstr)                % View the sparsity structure of Hstr
```



```
n = 1000;
xstart = -ones(n,1);
xstart(2:2:n,1) = 1;
options = optimoptions(@fminunc,'GradObj','on','HessPattern',Hstr);
[x,fval,exitflag,output] = fminunc(fun,xstart,options);
```

This 1000-variable problem is solved in seven iterations and seven conjugate gradient iterations with a positive `exitflag` indicating convergence. The final function value and measure of optimality at the solution `x` are both close to zero (for `fminunc`, the first-order optimality is the infinity norm of the gradient of the function, which is zero at a local minimum):

```
exitflag,fval,output
```

```
exitflag =  
    1  
  
fval =  
    7.4738e-17  
  
output =  
    iterations: 7  
    funcCount: 8  
    cgiterations: 7  
    firstorderopt: 7.9822e-10  
    algorithm: 'large-scale: trust-region Newton'  
    message: 'Local minimum found.'  
  
Optimization completed because the size of the grad...'  
    constrviolation: []
```

Constrained Nonlinear Optimization Algorithms

In this section...

“Constrained Optimization Definition” on page 6-26

“fmincon Trust Region Reflective Algorithm” on page 6-26

“fmincon Active Set Algorithm” on page 6-32

“fmincon SQP Algorithm” on page 6-42

“fmincon Interior Point Algorithm” on page 6-43

“fminbnd Algorithm” on page 6-47

“fseminf Problem Formulation and Algorithm” on page 6-47

Constrained Optimization Definition

Constrained minimization is the problem of finding a vector x that is a local minimum to a scalar function $f(x)$ subject to constraints on the allowable x :

$$\min_x f(x)$$

such that one or more of the following holds: $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, $l \leq x \leq u$. There are even more constraints used in semi-infinite programming; see “fseminf Problem Formulation and Algorithm” on page 6-47.

fmincon Trust Region Reflective Algorithm

Trust-Region Methods for Nonlinear Minimization

Many of the methods used in Optimization Toolbox solvers are based on *trust regions*, a simple yet powerful concept in optimization.

To understand the trust-region approach to optimization, consider the unconstrained minimization problem, minimize $f(x)$, where the function takes vector arguments and returns scalars. Suppose you are at a point x in n -space and you want to improve, i.e., move to a point with a lower function value. The basic idea is to approximate f with a simpler function q , which reasonably

reflects the behavior of function f in a neighborhood N around the point x . This neighborhood is the trust region. A trial step s is computed by minimizing (or approximately minimizing) over N . This is the trust-region subproblem,

$$\min_s \{q(s), s \in N\}. \quad (6-18)$$

The current point is updated to be $x + s$ if $f(x + s) < f(x)$; otherwise, the current point remains unchanged and N , the region of trust, is shrunk and the trial step computation is repeated.

The key questions in defining a specific trust-region approach to minimizing $f(x)$ are how to choose and compute the approximation q (defined at the current point x), how to choose and modify the trust region N , and how accurately to solve the trust-region subproblem. This section focuses on the unconstrained problem. Later sections discuss additional complications due to the presence of constraints on the variables.

In the standard trust-region method ([48]), the quadratic approximation q is defined by the first two terms of the Taylor approximation to F at x ; the neighborhood N is usually spherical or ellipsoidal in shape. Mathematically the trust-region subproblem is typically stated

$$\min \left\{ \frac{1}{2} s^T H s + s^T g \text{ such that } \|D s\| \leq \Delta \right\}, \quad (6-19)$$

where g is the gradient of f at the current point x , H is the Hessian matrix (the symmetric matrix of second derivatives), D is a diagonal scaling matrix, Δ is a positive scalar, and $\| \cdot \|$ is the 2-norm. Good algorithms exist for solving Equation 6-19 (see [48]); such algorithms typically involve the computation of a full eigensystem and a Newton process applied to the secular equation

$$\frac{1}{\Delta} - \frac{1}{\|s\|} = 0.$$

Such algorithms provide an accurate solution to Equation 6-19. However, they require time proportional to several factorizations of H . Therefore, for large-scale problems a different approach is needed. Several approximation and heuristic strategies, based on Equation 6-19, have been proposed in

the literature ([42] and [50]). The approximation approach followed in Optimization Toolbox solvers is to restrict the trust-region subproblem to a two-dimensional subspace S ([39] and [42]). Once the subspace S has been computed, the work to solve Equation 6-19 is trivial even if full eigenvalue/eigenvector information is needed (since in the subspace, the problem is only two-dimensional). The dominant work has now shifted to the determination of the subspace.

The two-dimensional subspace S is determined with the aid of a preconditioned conjugate gradient process described below. The solver defines S as the linear space spanned by s_1 and s_2 , where s_1 is in the direction of the gradient g , and s_2 is either an approximate Newton direction, i.e., a solution to

$$H \cdot s_2 = -g, \tag{6-20}$$

or a direction of negative curvature,

$$s_2^T \cdot H \cdot s_2 < 0. \tag{6-21}$$

The philosophy behind this choice of S is to force global convergence (via the steepest descent direction or negative curvature direction) and achieve fast local convergence (via the Newton step, when it exists).

A sketch of unconstrained minimization using trust-region ideas is now easy to give:

- 1** Formulate the two-dimensional trust-region subproblem.
- 2** Solve Equation 6-19 to determine the trial step s .
- 3** If $f(x + s) < f(x)$, then $x = x + s$.
- 4** Adjust Δ .

These four steps are repeated until convergence. The trust-region dimension Δ is adjusted according to standard rules. In particular, it is decreased if the trial step is not accepted, i.e., $f(x + s) \geq f(x)$. See [46] and [49] for a discussion of this aspect.

Optimization Toolbox solvers treat a few important special cases of f with specialized functions: nonlinear least-squares, quadratic functions, and linear least-squares. However, the underlying algorithmic ideas are the same as for the general case. These special cases are discussed in later sections.

Preconditioned Conjugate Gradient Method

A popular way to solve large symmetric positive definite systems of linear equations $Hp = -g$ is the method of Preconditioned Conjugate Gradients (PCG). This iterative approach requires the ability to calculate matrix-vector products of the form Hv where v is an arbitrary vector. The symmetric positive definite matrix M is a *preconditioner* for H . That is, $M = C^2$, where $C^{-1}HC^{-1}$ is a well-conditioned matrix or a matrix with clustered eigenvalues.

In a minimization context, you can assume that the Hessian matrix H is symmetric. However, H is guaranteed to be positive definite only in the neighborhood of a strong minimizer. Algorithm PCG exits when a direction of negative (or zero) curvature is encountered, i.e., $d^T H d \leq 0$. The PCG output direction, p , is either a direction of negative curvature or an approximate (*tol* controls how approximate) solution to the Newton system $Hp = -g$. In either case p is used to help define the two-dimensional subspace used in the trust-region approach discussed in “Trust-Region Methods for Nonlinear Minimization” on page 6-5.

Linear Equality Constraints

Linear constraints complicate the situation described for unconstrained minimization. However, the underlying ideas described previously can be carried through in a clean and efficient way. The trust-region methods in Optimization Toolbox solvers generate strictly feasible iterates.

The general linear equality constrained minimization problem can be written

$$\min\{f(x) \text{ such that } Ax = b\}, \quad (6-22)$$

where A is an m -by- n matrix ($m \leq n$). Some Optimization Toolbox solvers preprocess A to remove strict linear dependencies using a technique based on the LU factorization of A^T [46]. Here A is assumed to be of rank m .

The method used to solve Equation 6-22 differs from the unconstrained approach in two significant ways. First, an initial feasible point x_0 is computed, using a sparse least-squares step, so that $Ax_0 = b$. Second, Algorithm PCG is replaced with Reduced Preconditioned Conjugate Gradients (RPCG), see [46], in order to compute an approximate reduced Newton step (or a direction of negative curvature in the null space of A). The key linear algebra step involves solving systems of the form

$$\begin{bmatrix} C & \tilde{A}^T \\ \tilde{A} & 0 \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix}, \quad (6-23)$$

where \tilde{A} approximates A (small nonzeros of A are set to zero provided rank is not lost) and C is a sparse symmetric positive-definite approximation to H , i.e., $C = H$. See [46] for more details.

Box Constraints

The box constrained problem is of the form

$$\min \{f(x) \text{ such that } l \leq x \leq u\}, \quad (6-24)$$

where l is a vector of lower bounds, and u is a vector of upper bounds. Some (or all) of the components of l can be equal to $-\infty$ and some (or all) of the components of u can be equal to ∞ . The method generates a sequence of strictly feasible points. Two techniques are used to maintain feasibility while achieving robust convergence behavior. First, a scaled modified Newton step replaces the unconstrained Newton step (to define the two-dimensional subspace S). Second, reflections are used to increase the step size.

The scaled modified Newton step arises from examining the Kuhn-Tucker necessary conditions for Equation 6-24,

$$(D(x))^{-2} g = 0, \quad (6-25)$$

where

$$D(x) = \text{diag}(|v_k|^{-1/2}),$$

and the vector $v(x)$ is defined below, for each $1 \leq i \leq n$:

- If $g_i < 0$ and $u_i < \infty$ then $v_i = x_i - u_i$
- If $g_i \geq 0$ and $l_i > -\infty$ then $v_i = x_i - l_i$
- If $g_i < 0$ and $u_i = \infty$ then $v_i = -1$
- If $g_i \geq 0$ and $l_i = -\infty$ then $v_i = 1$

The nonlinear system Equation 6-25 is not differentiable everywhere. Nondifferentiability occurs when $v_i = 0$. You can avoid such points by maintaining strict feasibility, i.e., restricting $l < x < u$.

The scaled modified Newton step s_k for the nonlinear system of equations given by Equation 6-25 is defined as the solution to the linear system

$$\hat{M}Ds^N = -\hat{g} \quad (6-26)$$

at the k th iteration, where

$$\hat{g} = D^{-1}g = \text{diag}\left(|v|^{1/2}\right)g, \quad (6-27)$$

and

$$\hat{M} = D^{-1}HD^{-1} + \text{diag}(g)J^v. \quad (6-28)$$

Here J^v plays the role of the Jacobian of $|v|$. Each diagonal component of the diagonal matrix J^v equals 0, -1 , or 1 . If all the components of l and u are finite, $J^v = \text{diag}(\text{sign}(g))$. At a point where $g_i = 0$, v_i might not be differentiable.

$J_{ii}^v = 0$ is defined at such a point. Nondifferentiability of this type is not a cause for concern because, for such a component, it is not significant which value v_i takes. Further, $|v_i|$ will still be discontinuous at this point, but the function $|v_i|g_i$ is continuous.

Second, reflections are used to increase the step size. A (single) reflection step is defined as follows. Given a step p that intersects a bound constraint, consider the first bound constraint crossed by p ; assume it is the i th bound

constraint (either the i th upper or i th lower bound). Then the reflection step $p^R = p$ except in the i th component, where $p^R_i = -p_i$.

fmincon Active Set Algorithm

Introduction

In constrained optimization, the general aim is to transform the problem into an easier subproblem that can then be solved and used as the basis of an iterative process. A characteristic of a large class of early methods is the translation of the constrained problem to a basic unconstrained problem by using a penalty function for constraints that are near or beyond the constraint boundary. In this way the constrained problem is solved using a sequence of parameterized unconstrained optimizations, which in the limit (of the sequence) converge to the constrained problem. These methods are now considered relatively inefficient and have been replaced by methods that have focused on the solution of the Karush-Kuhn-Tucker (KKT) equations. The KKT equations are necessary conditions for optimality for a constrained optimization problem. If the problem is a so-called convex programming problem, that is, $f(x)$ and $G_i(x)$, $i = 1, \dots, m$, are convex functions, then the KKT equations are both necessary and sufficient for a global solution point.

Referring to GP (Equation 6-1), the Kuhn-Tucker equations can be stated as

$$\begin{aligned} \nabla f(x^*) + \sum_{i=1}^m \lambda_i \cdot \nabla G_i(x^*) &= 0 \\ \lambda_i \cdot G_i(x^*) &= 0, \quad i = 1, \dots, m_e \\ \lambda_i &\geq 0, \quad i = m_e + 1, \dots, m, \end{aligned} \tag{6-29}$$

in addition to the original constraints in Equation 6-1.

The first equation describes a canceling of the gradients between the objective function and the active constraints at the solution point. For the gradients to be canceled, Lagrange multipliers (λ_i , $i = 1, \dots, m$) are necessary to balance the deviations in magnitude of the objective function and constraint gradients. Because only active constraints are included in this canceling operation, constraints that are not active must not be included in this operation and so

are given Lagrange multipliers equal to 0. This is stated implicitly in the last two Kuhn-Tucker equations.

The solution of the KKT equations forms the basis to many nonlinear programming algorithms. These algorithms attempt to compute the Lagrange multipliers directly. Constrained quasi-Newton methods guarantee superlinear convergence by accumulating second-order information regarding the KKT equations using a quasi-Newton updating procedure. These methods are commonly referred to as Sequential Quadratic Programming (SQP) methods, since a QP subproblem is solved at each major iteration (also known as Iterative Quadratic Programming, Recursive Quadratic Programming, and Constrained Variable Metric methods).

The 'active-set' algorithm is not a large-scale algorithm; see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12.

Sequential Quadratic Programming (SQP)

SQP methods represent the state of the art in nonlinear programming methods. Schittkowski [36], for example, has implemented and tested a version that outperforms every other tested method in terms of efficiency, accuracy, and percentage of successful solutions, over a large number of test problems.

Based on the work of Biggs [1], Han [22], and Powell ([32] and [33]), the method allows you to closely mimic Newton's method for constrained optimization just as is done for unconstrained optimization. At each major iteration, an approximation is made of the Hessian of the Lagrangian function using a quasi-Newton updating method. This is then used to generate a QP subproblem whose solution is used to form a search direction for a line search procedure. An overview of SQP is found in Fletcher [13], Gill et al. [19], Powell [35], and Schittkowski [23]. The general method, however, is stated here.

Given the problem description in GP (Equation 6-1) the principal idea is the formulation of a QP subproblem based on a quadratic approximation of the Lagrangian function.

$$L(x, \lambda) = f(x) + \sum_{i=1}^m \lambda_i \cdot g_i(x). \quad (6-30)$$

Here you simplify Equation 6-1 by assuming that bound constraints have been expressed as inequality constraints. You obtain the QP subproblem by linearizing the nonlinear constraints.

Quadratic Programming (QP) Subproblem

$$\begin{aligned} \min_{d \in \mathcal{R}^n} & \frac{1}{2} d^T H_k d + \nabla f(x_k)^T d \\ & \nabla g_i(x_k)^T d + g_i(x_k) = 0, \quad i = 1, \dots, m_e \\ & \nabla g_i(x_k)^T d + g_i(x_k) \leq 0, \quad i = m_e + 1, \dots, m. \end{aligned} \quad (6-31)$$

This subproblem can be solved using any QP algorithm (see, for instance, “Quadratic Programming Solution” on page 6-37). The solution is used to form a new iterate

$$x_{k+1} = x_k + \alpha_k d_k.$$

The step length parameter α_k is determined by an appropriate line search procedure so that a sufficient decrease in a merit function is obtained (see “Updating the Hessian Matrix” on page 6-35). The matrix H_k is a positive definite approximation of the Hessian matrix of the Lagrangian function (Equation 6-30). H_k can be updated by any of the quasi-Newton methods, although the BFGS method (see “Updating the Hessian Matrix” on page 6-35) appears to be the most popular.

A nonlinearly constrained problem can often be solved in fewer iterations than an unconstrained problem using SQP. One of the reasons for this is that, because of limits on the feasible area, the optimizer can make informed decisions regarding directions of search and step length.

Consider Rosenbrock’s function with an additional nonlinear inequality constraint, $g(x)$,

$$x_1^2 + x_2^2 - 1.5 \leq 0. \quad (6-32)$$

This was solved by an SQP implementation in 96 iterations compared to 140 for the unconstrained case. SQP Method on Nonlinearly Constrained Rosenbrock's Function (Equation 6-6) on page 6-35 shows the path to the solution point $x = [0.9072, 0.8228]$ starting at $x = [-1.9, 2.0]$.




Figure 6-3: SQP Method on Nonlinearly Constrained Rosenbrock's Function (Equation 6-6)

SQP Implementation

The SQP implementation consists of three main stages, which are discussed briefly in the following subsections:

- “Updating the Hessian Matrix” on page 6-35
- “Quadratic Programming Solution” on page 6-37
- “Line Search and Merit Function” on page 6-41

Updating the Hessian Matrix. At each major iteration a positive definite quasi-Newton approximation of the Hessian of the Lagrangian function, H , is calculated using the BFGS method, where λ_i , $i = 1, \dots, m$, is an estimate of the Lagrange multipliers.

$$H_{k+1} = H_k + \frac{q_k q_k^T}{q_k^T s_k} - \frac{H_k s_k s_k^T H_k^T}{s_k^T H_k s_k}, \quad (6-33)$$

where

$$s_k = x_{k+1} - x_k$$

$$q_k = \left(\nabla f(x_{k+1}) + \sum_{i=1}^m \lambda_i \cdot \nabla g_i(x_{k+1}) \right) - \left(\nabla f(x_k) + \sum_{i=1}^m \lambda_i \cdot \nabla g_i(x_k) \right).$$

Powell [33] recommends keeping the Hessian positive definite even though it might be positive indefinite at the solution point. A positive definite

Hessian is maintained providing $q_k^T s_k$ is positive at each update and that H is initialized with a positive definite matrix. When $q_k^T s_k$ is not positive, q_k

is modified on an element-by-element basis so that $q_k^T s_k > 0$. The general aim of this modification is to distort the elements of q_k , which contribute to a positive definite update, as little as possible. Therefore, in the initial phase of the modification, the most negative element of $q_k^* s_k$ is repeatedly halved.

This procedure is continued until $q_k^T s_k$ is greater than or equal to a small negative tolerance. If, after this procedure, $q_k^T s_k$ is still not positive, modify q_k by adding a vector v multiplied by a constant scalar w , that is,

$$q_k = q_k + wv, \quad (6-34)$$

where

$$v_i = \nabla g_i(x_{k+1}) \cdot g_i(x_{k+1}) - \nabla g_i(x_k) \cdot g_i(x_k)$$

$$\text{if } (q_k)_i \cdot w < 0 \text{ and } (q_k)_i \cdot (s_k)_i < 0, i = 1, \dots, m$$

$$v_i = 0 \text{ otherwise,}$$

and increase w systematically until $q_k^T s_k$ becomes positive.

The functions `fmincon`, `fminimax`, `fgoalattain`, and `fseminf` all use SQP. If `Display` is set to 'iter' in options, then various information is given

such as function values and the maximum constraint violation. When the Hessian has to be modified using the first phase of the preceding procedure to keep it positive definite, then `Hessian modified` is displayed. If the Hessian has to be modified again using the second phase of the approach described above, then `Hessian modified twice` is displayed. When the QP subproblem is infeasible, then `infeasible` is displayed. Such displays are usually not a cause for concern but indicate that the problem is highly nonlinear and that convergence might take longer than usual. Sometimes the message `no update` is displayed, indicating that $q_k^T s_k$ is nearly zero. This can be an indication that the problem setup is wrong or you are trying to minimize a noncontinuous function.

Quadratic Programming Solution. At each major iteration of the SQP method, a QP problem of the following form is solved, where A_i refers to the i th row of the m -by- n matrix A .

$$\begin{aligned} \min_{d \in \mathcal{R}^n} q(d) &= \frac{1}{2} d^T H d + c^T d, \\ A_i d &= b_i, \quad i = 1, \dots, m_e \\ A_i d &\leq b_i, \quad i = m_e + 1, \dots, m. \end{aligned} \tag{6-35}$$

The method used in Optimization Toolbox functions is an active set strategy (also known as a projection method) similar to that of Gill et al., described in [18] and [17]. It has been modified for both Linear Programming (LP) and Quadratic Programming (QP) problems.

The solution procedure involves two phases. The first phase involves the calculation of a feasible point (if one exists). The second phase involves the generation of an iterative sequence of feasible points that converge to the solution. In this method an active set, \bar{A}_k , is maintained that is an estimate of the active constraints (i.e., those that are on the constraint boundaries) at the solution point. Virtually all QP algorithms are active set methods. This point is emphasized because there exist many different methods that are very similar in structure but that are described in widely different terms.

\bar{A}_k is updated at each iteration k , and this is used to form a basis for a search direction \hat{d}_k . Equality constraints always remain in the active set \bar{A}_k . The

notation for the variable \hat{d}_k is used here to distinguish it from d_k in the major iterations of the SQP method. The search direction \hat{d}_k is calculated and minimizes the objective function while remaining on any active constraint boundaries. The feasible subspace for \hat{d}_k is formed from a basis Z_k whose columns are orthogonal to the estimate of the active set \bar{A}_k (i.e., $\bar{A}_k Z_k = 0$). Thus a search direction, which is formed from a linear summation of any combination of the columns of Z_k , is guaranteed to remain on the boundaries of the active constraints.

The matrix Z_k is formed from the last $m - l$ columns of the QR decomposition of the matrix \bar{A}_k^T , where l is the number of active constraints and $l < m$. That is, Z_k is given by

$$Z_k = Q[:, l+1 : m], \quad (6-36)$$

where

$$Q^T \bar{A}_k^T = \begin{bmatrix} R \\ 0 \end{bmatrix}.$$

Once Z_k is found, a new search direction \hat{d}_k is sought that minimizes $q(d)$ where \hat{d}_k is in the null space of the active constraints. That is, \hat{d}_k is a linear combination of the columns of Z_k : $\hat{d}_k = Z_k p$ for some vector p .

Then if you view the quadratic as a function of p , by substituting for \hat{d}_k , you have

$$q(p) = \frac{1}{2} p^T Z_k^T H Z_k p + c^T Z_k p. \quad (6-37)$$

Differentiating this with respect to p yields

$$\nabla q(p) = Z_k^T H Z_k p + Z_k^T c. \quad (6-38)$$

$\nabla q(p)$ is referred to as the projected gradient of the quadratic function because it is the gradient projected in the subspace defined by Z_k . The term $Z_k^T H Z_k$ is called the projected Hessian. Assuming the Hessian matrix H is positive definite (which is the case in this implementation of SQP), then the minimum of the function $q(p)$ in the subspace defined by Z_k occurs when $\nabla q(p) = 0$, which is the solution of the system of linear equations

$$Z_k^T H Z_k p = -Z_k^T c. \quad (6-39)$$

A step is then taken of the form

$$x_{k+1} = x_k + \alpha \hat{d}_k, \quad \text{where } \hat{d}_k = Z_k^T p. \quad (6-40)$$

At each iteration, because of the quadratic nature of the objective function, there are only two choices of step length α . A step of unity along \hat{d}_k is the exact step to the minimum of the function restricted to the null space of \bar{A}_k . If such a step can be taken, without violation of the constraints, then this is the solution to QP (Equation 6-35). Otherwise, the step along \hat{d}_k to the nearest constraint is less than unity and a new constraint is included in the active set at the next iteration. The distance to the constraint boundaries in any direction \hat{d}_k is given by

$$\alpha = \min_{i \in \{1, \dots, m\}} \left\{ \frac{-(A_i x_k - b_i)}{A_i \hat{d}_k} \right\}, \quad (6-41)$$

which is defined for constraints not in the active set, and where the direction \hat{d}_k is towards the constraint boundary, i.e., $A_i \hat{d}_k > 0$, $i = 1, \dots, m$.

When n independent constraints are included in the active set, without location of the minimum, Lagrange multipliers, λ_k , are calculated that satisfy the nonsingular set of linear equations

$$\bar{A}_k^T \lambda_k = c. \quad (6-42)$$

If all elements of λ_k are positive, x_k is the optimal solution of QP (Equation 6-35). However, if any component of λ_k is negative, and the component does not correspond to an equality constraint, then the corresponding element is deleted from the active set and a new iterate is sought.

Initialization

The algorithm requires a feasible point to start. If the current point from the SQP method is not feasible, then you can find a point by solving the linear programming problem

$$\begin{aligned} \min_{\gamma \in \mathfrak{R}, x \in \mathfrak{R}^n} \quad & \gamma \text{ such that} \\ A_i x = b_i, \quad & i = 1, \dots, m_e \\ A_i x - \gamma \leq b_i, \quad & i = m_e + 1, \dots, m. \end{aligned} \quad (6-43)$$

The notation A_i indicates the i th row of the matrix A . You can find a feasible point (if one exists) to Equation 6-43 by setting x to a value that satisfies the equality constraints. You can determine this value by solving an under- or overdetermined set of linear equations formed from the set of equality constraints. If there is a solution to this problem, then the slack variable γ is set to the maximum inequality constraint at this point.

You can modify the preceding QP algorithm for LP problems by setting the search direction to the steepest descent direction at each iteration, where g_k is the gradient of the objective function (equal to the coefficients of the linear objective function).

$$\hat{d}_k = -Z_k Z_k^T g_k. \quad (6-44)$$

If a feasible point is found using the preceding LP method, the main QP phase is entered. The search direction \hat{d}_k is initialized with a search direction \hat{d}_1 found from solving the set of linear equations

$$H\hat{d}_1 = -g_k, \quad (6-45)$$

where g_k is the gradient of the objective function at the current iterate x_k (i.e., $Hx_k + c$).

If a feasible solution is not found for the QP problem, the direction of search for the main SQP routine \hat{d}_k is taken as one that minimizes γ .

Line Search and Merit Function. The solution to the QP subproblem produces a vector d_k , which is used to form a new iterate

$$x_{k+1} = x_k + \alpha d_k. \quad (6-46)$$

The step length parameter α_k is determined in order to produce a sufficient decrease in a merit function. The merit function used by Han [22] and Powell [33] of the following form is used in this implementation.

$$\Psi(x) = f(x) + \sum_{i=1}^{m_e} r_i \cdot g_i(x) + \sum_{i=m_e+1}^m r_i \cdot \max[0, g_i(x)]. \quad (6-47)$$

Powell recommends setting the penalty parameter

$$r_i = (r_{k+1})_i = \max_i \left\{ \lambda_i, \frac{(r_k)_i + \lambda_i}{2} \right\}, \quad i = 1, \dots, m. \quad (6-48)$$

This allows positive contribution from constraints that are inactive in the QP solution but were recently active. In this implementation, the penalty parameter r_i is initially set to

$$r_i = \frac{\|\nabla f(x)\|}{\|\nabla g_i(x)\|}, \quad (6-49)$$

where $\| \cdot \|$ represents the Euclidean norm.

This ensures larger contributions to the penalty parameter from constraints with smaller gradients, which would be the case for active constraints at the solution point.

fmincon SQP Algorithm

The sqp algorithm is similar to the active-set algorithm (for a description, see “fmincon Active Set Algorithm” on page 6-32). The basic sqp algorithm is described in Chapter 18 of Nocedal and Wright [31].

The most important differences between the sqp and the active-set algorithms are:

Strict Feasibility With Respect to Bounds

The sqp algorithm takes every iterative step in the region constrained by bounds. Furthermore, finite difference steps also respect bounds. Bounds are not strict; a step can be exactly on a boundary. This strict feasibility can be beneficial when your objective function or nonlinear constraint functions are undefined or are complex outside the region constrained by bounds.

Robustness to Non-Double Results

During its iterations, the sqp algorithm can attempt to take a step that fails. This means an objective function or nonlinear constraint function you supply returns a value of Inf, NaN, or a complex value. In this case, the algorithm attempts to take a smaller step.

Refactored Linear Algebra Routines

The sqp algorithm uses a different set of linear algebra routines to solve the quadratic programming subproblem, Equation 6-31. These routines are more efficient in both memory usage and speed than the active-set routines.

Reformulated Feasibility Routines

The sqp algorithm has two new approaches to the solution of Equation 6-31 when constraints are not satisfied.

- The sqp algorithm combines the objective and constraint functions into a merit function. The algorithm attempts to minimize the merit function

subject to relaxed constraints. This modified problem can lead to a feasible solution. However, this approach has more variables than the original problem, so the problem size in Equation 6-31 increases. The increased size can slow the solution of the subproblem. These routines are based on the articles by Spellucci [60] and Tone [61]. The sqp algorithm sets the penalty parameter for the merit function Equation 6-47 according to the suggestion in [41].

- Suppose nonlinear constraints are not satisfied, and an attempted step causes the constraint violation to grow. The sqp algorithm attempts to obtain feasibility using a second-order approximation to the constraints. The second-order technique can lead to a feasible solution. However, this technique can slow the solution by requiring more evaluations of the nonlinear constraint functions.

fmincon Interior Point Algorithm

Barrier Function

The interior-point approach to constrained minimization is to solve a sequence of approximate minimization problems. The original problem is

$$\min_x f(x), \text{ subject to } h(x) = 0 \text{ and } g(x) \leq 0. \quad (6-50)$$

For each $\mu > 0$, the approximate problem is

$$\min_{x,s} f_\mu(x,s) = \min_{x,s} f(x) - \mu \sum_i \ln(s_i), \text{ subject to } h(x) = 0 \text{ and } g(x) + s = 0. \quad (6-51)$$

There are as many slack variables s_i as there are inequality constraints g . The s_i are restricted to be positive to keep $\ln(s_i)$ bounded. As μ decreases to zero, the minimum of f_μ should approach the minimum of f . The added logarithmic term is called a *barrier function*. This method is described in [40], [41], and [51].

The approximate problem Equation 6-51 is a sequence of equality constrained problems. These are easier to solve than the original inequality-constrained problem Equation 6-50.

To solve the approximate problem, the algorithm uses one of two main types of steps at each iteration:

- A *direct* step in (x, s) . This step attempts to solve the KKT equations, Equation 3-2 and Equation 3-3, for the approximate problem via a linear approximation. This is also called a *Newton step*.
- A *CG* (conjugate gradient) step, using a trust region.

By default, the algorithm first attempts to take a direct step. If it cannot, it attempts a CG step. One case where it does not take a direct step is when the approximate problem is not locally convex near the current iterate.

At each iteration the algorithm decreases a *merit function*, such as

$$f_{\mu}(x, s) + \nu \|(h(x), g(x) + s)\|.$$

The parameter ν may increase with iteration number in order to force the solution towards feasibility. If an attempted step does not decrease the merit function, the algorithm rejects the attempted step, and attempts a new step.

If either the objective function or a nonlinear constraint function returns a complex value, NaN, Inf, or an error at an iterate x_j , the algorithm rejects x_j . The rejection has the same effect as if the merit function did not decrease sufficiently: the algorithm then attempts a different, shorter step. Wrap any code that can error in try-catch:

```
function val = userFcn(x)
try
    val = ... % code that can error
catch
    val = NaN;
end
```

The objective and constraints must yield proper (double) values at the initial point.

Direct Step

The following variables are used in defining the direct step:

- H denotes the Hessian of the Lagrangian of f_μ :

$$H = \nabla^2 f(x) + \sum_i \lambda_i \nabla^2 g_i(x) + \sum_j \lambda_j \nabla^2 h_j(x). \quad (6-52)$$

- J_g denotes the Jacobian of the constraint function g .
- J_h denotes the Jacobian of the constraint function h .
- $S = \text{diag}(s)$.
- λ denotes the Lagrange multiplier vector associated with constraints g
- $\Lambda = \text{diag}(\lambda)$.
- y denotes the Lagrange multiplier vector associated with h .
- e denote the vector of ones the same size as g .

Equation 6-53 defines the direct step $(\Delta x, \Delta s)$:

$$\begin{bmatrix} H & 0 & J_h^T & J_g^T \\ 0 & S\Lambda & 0 & -S \\ J_h & 0 & I & 0 \\ J_g & -S & 0 & I \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta s \\ -\Delta y \\ -\Delta \lambda \end{bmatrix} = - \begin{bmatrix} \nabla f - J_h^T y - J_g^T \lambda \\ S\lambda - \mu e \\ h \\ g + s \end{bmatrix}. \quad (6-53)$$

This equation comes directly from attempting to solve Equation 3-2 and Equation 3-3 using a linearized Lagrangian.

In order to solve this equation for $(\Delta x, \Delta s)$, the algorithm makes an LDL factorization of the matrix. (See Example 3 — The Structure of D in the MATLAB `ldl` function reference page.) This is the most computationally expensive step. One result of this factorization is a determination of whether the projected Hessian is positive definite or not; if not, the algorithm uses a conjugate gradient step, described in the next section.

Conjugate Gradient Step

The conjugate gradient approach to solving the approximate problem Equation 6-51 is similar to other conjugate gradient calculations. In this case, the algorithm adjusts both x and s , keeping the slacks s positive. The

approach is to minimize a quadratic approximation to the approximate problem in a trust region, subject to linearized constraints.

Specifically, let R denote the radius of the trust region, and let other variables be defined as in “Direct Step” on page 6-44. The algorithm obtains Lagrange multipliers by approximately solving the KKT equations

$$\nabla_x L = \nabla_x f(x) + \sum_i \lambda_i \nabla g_i(x) + \sum_j y_j \nabla h_j(x) = 0,$$

in the least-squares sense, subject to λ being positive. Then it takes a step $(\Delta x, \Delta s)$ to approximately solve

$$\min_{\Delta x, \Delta s} \nabla f^T \Delta x + \frac{1}{2} \Delta x^T \nabla_{xx}^2 L \Delta x + \mu e^T S^{-1} \Delta s + \frac{1}{2} \Delta s^T S^{-1} \Lambda \Delta s, \quad (6-54)$$

subject to the linearized constraints

$$g(x) + J_g \Delta x + \Delta s = 0, \quad h(x) + J_h \Delta x = 0. \quad (6-55)$$

To solve Equation 6-55, the algorithm tries to minimize a norm of the linearized constraints inside a region with radius scaled by R . Then Equation 6-54 is solved with the constraints being to match the residual from solving Equation 6-55, staying within the trust region of radius R , and keeping s strictly positive. For details of the algorithm and the derivation, see [40], [41], and [51]. For another description of conjugate gradients, see “Preconditioned Conjugate Gradient Method” on page 6-29.

Interior-Point Algorithm Options

Here are the meanings and effects of several options in the interior-point algorithm.

- `AlwaysHonorConstraints` — When set to 'bounds', every iterate satisfies the bound constraints you have set. When set to 'none', the algorithm may violate bounds during intermediate iterations.
- `Hessian` — When set to:
 - 'user-supplied', pass the Hessian of the Lagrangian in a user-supplied function, whose function handle must be given in the option `HessFcn`.

- 'bfgs', `fmincon` calculates the Hessian by a dense quasi-Newton approximation.
- 'lbfgs', `fmincon` calculates the Hessian by a limited-memory, large-scale quasi-Newton approximation.
- 'fin-diff-grads', `fmincon` calculates a Hessian-times-vector product by finite differences of the gradient(s); other options need to be set appropriately.

You can also give a separate function for Hessian-times-vector. See “Hessian” on page 10-49 for more details on these options.

- `InitBarrierParam` — The starting value for μ . By default, this is 0.1.
- `ScaleProblem` — When set to 'obj-and-constr', the algorithm works with scaled versions of the objective function and constraints. It carefully scales them by their initial values. Disable scaling by setting `ScaleProblem` to 'none'.
- `SubproblemAlgorithm` — Determines whether or not to attempt the direct Newton step. The default setting 'ldl-factorization' allows this type of step to be attempted. The setting 'cg' allows only conjugate gradient steps.

For a complete list of options see “Interior-Point Algorithm” on page 10-60.

fminbnd Algorithm

`fminbnd` is a solver available in any MATLAB installation. It solves for a local minimum in one dimension within a bounded interval. It is not based on derivatives. Instead, it uses golden-section search and parabolic interpolation.

fseminf Problem Formulation and Algorithm

fseminf Problem Formulation

`fseminf` addresses optimization problems with additional types of constraints compared to those addressed by `fmincon`. The formulation of `fmincon` is

$$\min_x f(x)$$

such that $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, and $l \leq x \leq u$.

`fseminf` adds the following set of semi-infinite constraints to those already given. For w_j in a one- or two-dimensional bounded interval or rectangle I_j , for a vector of continuous functions $K(x, w)$, the constraints are

$$K_j(x, w_j) \leq 0 \text{ for all } w_j \in I_j.$$

The term “dimension” of an `fseminf` problem means the maximal dimension of the constraint set I : 1 if all I_j are intervals, and 2 if at least one I_j is a rectangle. The size of the vector of K does not enter into this concept of dimension.

The reason this is called semi-infinite programming is that there are a finite number of variables (x and w_j), but an infinite number of constraints. This is because the constraints on x are over a set of continuous intervals or rectangles I_j , which contains an infinite number of points, so there are an infinite number of constraints: $K_j(x, w_j) \leq 0$ for an infinite number of points w_j .

You might think a problem with an infinite number of constraints is impossible to solve. `fseminf` addresses this by reformulating the problem to an equivalent one that has two stages: a maximization and a minimization. The semi-infinite constraints are reformulated as

$$\max_{w_j \in I_j} K_j(x, w_j) \leq 0 \text{ for all } j = 1, \dots, |K|, \quad (6-56)$$

where $|K|$ is the number of components of the vector K ; i.e., the number of semi-infinite constraint functions. For fixed x , this is an ordinary maximization over bounded intervals or rectangles.

`fseminf` further simplifies the problem by making piecewise quadratic or cubic approximations $\kappa_j(x, w_j)$ to the functions $K_j(x, w_j)$, for each x that the solver visits. `fseminf` considers only the maxima of the interpolation function $\kappa_j(x, w_j)$, instead of $K_j(x, w_j)$, in Equation 6-56. This reduces the original problem, minimizing a semi-infinitely constrained function, to a problem with a finite number of constraints.

Sampling Points. Your semi-infinite constraint function must provide a set of sampling points, points used in making the quadratic or cubic approximations. To accomplish this, it should contain:

- The initial spacing s between sampling points w
- A way of generating the set of sampling points w from s

The initial spacing s is a $|K|$ -by-2 matrix. The j th row of s represents the spacing for neighboring sampling points for the constraint function K_j . If K_j depends on a one-dimensional w_j , set $s(j,2) = 0$. `fseminf` updates the matrix s in subsequent iterations.

`fseminf` uses the matrix s to generate the sampling points w , which it then uses to create the approximation $\kappa_j(x, w_j)$. Your procedure for generating w from s should keep the same intervals or rectangles I_j during the optimization.

Example of Creating Sampling Points. Consider a problem with two semi-infinite constraints, K_1 and K_2 . K_1 has one-dimensional w_1 , and K_2 has two-dimensional w_2 . The following code generates a sampling set from $w_1 = 2$ to 12:

```
% Initial sampling interval
if isnan(s(1,1))
    s(1,1) = .2;
    s(1,2) = 0;
end

% Sampling set
w1 = 2:s(1,1):12;
```

`fseminf` specifies s as NaN when it first calls your constraint function. Checking for this allows you to set the initial sampling interval.

The following code generates a sampling set from w_2 in a square, with each component going from 1 to 100, initially sampled more often in the first component than the second:

```
% Initial sampling interval
if isnan(s(1,1))
    s(2,1) = 0.2;
    s(2,2) = 0.5;
end

% Sampling set
```

```
w2x = 1:s(2,1):100;
w2y = 1:s(2,2):100;
[wx,wy] = meshgrid(w2x,w2y);
```

The preceding code snippets can be simplified as follows:

```
% Initial sampling interval
if isnan(s(1,1))
    s = [0.2 0;0.2 0.5];
end

% Sampling set
w1 = 2:s(1,1):12;
w2x = 1:s(2,1):100;
w2y = 1:s(2,2):100;
[wx,wy] = meshgrid(w2x,w2y);
```

fseminf Algorithm

fseminf essentially reduces the problem of semi-infinite programming to a problem addressed by fmincon. fseminf takes the following steps to solve semi-infinite programming problems:

- 1 At the current value of x , fseminf identifies all the $w_{j,i}$ such that the interpolation $\kappa_j(x, w_{j,i})$ is a local maximum. (The maximum refers to varying w for fixed x .)
- 2 fseminf takes one iteration step in the solution of the fmincon problem:

$$\min_x f(x)$$

such that $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, and $l \leq x \leq u$, where $c(x)$ is augmented with all the maxima of $\kappa_j(x, w_j)$ taken over all $w_j I_j$, which is equal to the maxima over j and i of $\kappa_j(x, w_{j,i})$.

- 3 fseminf checks if any stopping criterion is met at the new point x (to halt the iterations); if not, it continues to step 4.

- 4** `fseminf` checks if the discretization of the semi-infinite constraints needs updating, and updates the sampling points appropriately. This provides an updated approximation $\kappa_j(x, w_j)$. Then it continues at step 1.

Nonlinear Inequality Constraints

This example shows how to solve a scalar minimization problem with nonlinear inequality constraints. The problem is to find x that solves

$$\min_x f(x) = e^{x_1} (4x_1^2 + 2x_2^2 + 4x_1x_2 + 2x_2 + 1). \quad (6-57)$$

subject to the constraints

$$\begin{aligned} x_1x_2 - x_1 - x_2 &\leq -1.5, \\ x_1x_2 &\geq -10. \end{aligned}$$

Because neither of the constraints is linear, you cannot pass the constraints to `fmincon` at the command line. Instead you can create a second file, `confun.m`, that returns the value at both constraints at the current x in a vector c . The constrained optimizer, `fmincon`, is then invoked. Because `fmincon` expects the constraints to be written in the form $c(x) \leq 0$, you must rewrite your constraints in the form

$$\begin{aligned} x_1x_2 - x_1 - x_2 + 1.5 &\leq 0, \\ -x_1x_2 - 10 &\leq 0. \end{aligned} \quad (6-58)$$

Step 1: Write a file `objfun.m` for the objective function.

```
function f = objfun(x)
f = exp(x(1))*(4*x(1)^2 + 2*x(2)^2 + 4*x(1)*x(2) + 2*x(2) + 1);
```

Step 2: Write a file `confun.m` for the constraints.

```
function [c, ceq] = confun(x)
% Nonlinear inequality constraints
c = [1.5 + x(1)*x(2) - x(1) - x(2);
     -x(1)*x(2) - 10];
% Nonlinear equality constraints
ceq = [];
```

Step 3: Invoke constrained optimization routine.

```
x0 = [-1,1];      % Make a starting guess at the solution
options = optimoptions(@fmincon,'Algorithm','sqp');
[x,fval] = ...
fmincon(@objfun,x0,[],[],[],[],[],[],[],@confun,options);
```

fmincon produces the solution x with function value fval:

```
x, fval
x =
   -9.5474    1.0474
fval =
    0.0236
```

You can evaluate the constraints at the solution by entering

```
[c,ceq] = confun(x)
```

This returns numbers close to zero, such as

```
c =

    1.0e-14 *
   -0.6661
    0.7105
```

```
ceq =

    []
```

Note that both constraint values are, to within a small tolerance, less than or equal to 0; that is, x satisfies $c(x) \leq 0$.

Nonlinear Constraints with Gradients

Ordinarily the medium-scale minimization routines use numerical gradients calculated by finite-difference approximation. This procedure systematically perturbs each of the variables in order to calculate function and constraint partial derivatives. Alternatively, you can provide a function to compute partial derivatives analytically. Typically, the problem is solved more accurately and efficiently if such a function is provided.

To solve

$$\min_x f(x) = e^{x_1} (4x_1^2 + 2x_2^2 + 4x_1x_2 + 2x_2 + 1).$$

using analytically determined gradients, do the following.

Step 1: Write a file for the objective function and gradient.

```
function [f,gradf] = objfungrad(x)
f = exp(x(1))*(4*x(1)^2+2*x(2)^2+4*x(1)*x(2)+2*x(2)+1);
% Gradient of the objective function
if nargin > 1
    gradf = [ f + exp(x(1)) * (8*x(1) + 4*x(2)),
             exp(x(1))*(4*x(1)+4*x(2)+2)];
end
```

Step 2: Write a file for the nonlinear constraints and the gradients of the nonlinear constraints.

```
function [c,ceq,DC,DCEq] = confungrad(x)
c(1) = 1.5 + x(1) * x(2) - x(1) - x(2); %Inequality constraints
c(2) = -x(1) * x(2)-10;
% No nonlinear equality constraints
ceq=[];
% Gradient of the constraints
if nargin > 2
    DC= [x(2)-1, -x(2);
         x(1)-1, -x(1)];
```



```

DCeq = [];
end

```

`gradf` contains the partial derivatives of the objective function, f , returned by `objfungrad(x)`, with respect to each of the elements in x :

$$\nabla f = \begin{bmatrix} e^{x_1} (4x_1^2 + 2x_2^2 + 4x_1x_2 + 2x_2 + 1) + e^{x_1} (8x_1 + 4x_2) \\ e^{x_1} (4x_1 + 4x_2 + 2) \end{bmatrix}. \quad (6-59)$$

The columns of `DC` contain the partial derivatives for each respective constraint (i.e., the i th column of `DC` is the partial derivative of the i th constraint with respect to x). So in the above example, `DC` is

$$\begin{bmatrix} \frac{\partial c_1}{\partial x_1} & \frac{\partial c_2}{\partial x_1} \\ \frac{\partial c_1}{\partial x_2} & \frac{\partial c_2}{\partial x_2} \end{bmatrix} = \begin{bmatrix} x_2 - 1 & -x_2 \\ x_1 - 1 & -x_1 \end{bmatrix}. \quad (6-60)$$

Since you are providing the gradient of the objective in `objfungrad.m` and the gradient of the constraints in `confungrad.m`, you *must* tell `fmincon` that these files contain this additional information. Use `optimoptions` to turn the options `GradObj` and `GradConstr` to 'on' in the example's existing options:

```
options = optimoptions(options,'GradObj','on','GradConstr','on');
```

If you do not set these options to 'on' in the options structure, `fmincon` does not use the analytic gradients.

The arguments `lb` and `ub` place lower and upper bounds on the independent variables in x . In this example, there are no bound constraints and so they are both set to `[]`.

Step 3: Invoke the constrained optimization routine.

```

x0 = [-1,1];           % Starting guess
options = optimoptions(@fmincon,'Algorithm','sqp');
options = optimoptions(options,'GradObj','on','GradConstr','on');
lb = []; ub = [];     % No upper or lower bounds

```

```
[x,fval] = fmincon(@objfungrad,x0,[],[],[],[],lb,ub,...
    @confungrad,options);
```

The results:

```
x,fval
```

```
x =
    -9.5474    1.0474
```

```
fval =
    0.0236
```

```
[c,ceq] = confungrad(x) % Check the constraint values at x
```

```
c =
    1.0e-13 *
    -0.1066
    0.1066
```

```
ceq =
    []
```

fmincon Interior-Point Algorithm with Analytic Hessian

The `fmincon` interior-point algorithm can accept a Hessian function as an input. When you supply a Hessian, you may obtain a faster, more accurate solution to a constrained minimization problem.

The constraint set for this example is the intersection of the interior of two cones—one pointing up, and one pointing down. The constraint function `c` is a two-component vector, one component for each cone. Since this is a three-dimensional example, the gradient of the constraint `c` is a 3-by-2 matrix.

```
function [c ceq gradc gradceq] = twocone(x)
% This constraint is two cones, z > -10 + r
% and z < 3 - r

ceq = [];
r = sqrt(x(1)^2 + x(2)^2);
c = [-10+r-x(3);
     x(3)-3+r];

if nargin > 2

    gradceq = [];
    gradc = [x(1)/r,x(1)/r;
            x(2)/r,x(2)/r;
            -1,1];

end
```


The objective function grows rapidly negative as the `x(1)` coordinate becomes negative. Its gradient is a three-element vector.

```
function [f gradf] = bigtopleft(x)
% This is a simple function that grows rapidly negative
% as x(1) gets negative
%
f=10*x(1)^3+x(1)*x(2)^2+x(3)*(x(1)^2+x(2)^2);

if nargin > 1
```

```
gradf=[30*x(1)^2+x(2)^2+2*x(3)*x(1);  
       2*x(1)*x(2)+2*x(3)*x(2);  
       (x(1)^2+x(2)^2)];  
  
end
```

Here is a plot of the problem. The shading represents the value of the objective function. You can see that the objective function is minimized near $x = [-6.5, 0, -3.5]$:



The Hessian of the Lagrangian is given by the equation:

$$\nabla_{xx}^2 L(x, \lambda) = \nabla^2 f(x) + \sum \lambda_i \nabla^2 c_i(x) + \sum \lambda_i \nabla^2 ceq_i(x).$$

The following function computes the Hessian at a point x with Lagrange multiplier structure `lambda`:

```
function h = hessinterior(x,lambda)
```

```

h = [60*x(1)+2*x(3),2*x(2),2*x(1);
     2*x(2),2*(x(1)+x(3)),2*x(2);
     2*x(1),2*x(2),0];% Hessian of f
r = sqrt(x(1)^2+x(2)^2);% radius
rinv3 = 1/r^3;
hessc = [(x(2))^2*rinv3,-x(1)*x(2)*rinv3,0;
         -x(1)*x(2)*rinv3,x(1)^2*rinv3,0;
         0,0,0];% Hessian of both c(1) and c(2)
h = h + lambda.ineqnonlin(1)*hessc + lambda.ineqnonlin(2)*hessc;

```

Run this problem using the interior-point algorithm in `fmincon`. To do this using the Optimization app:

- 1 Set the problem as in the following figure.

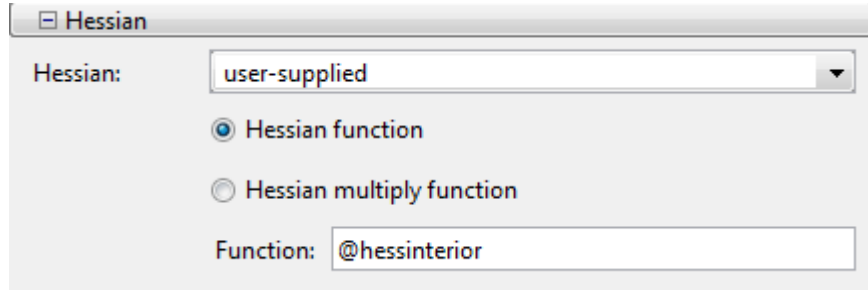
The screenshot shows the Optimization app interface with the following settings:

- Solver:** `fmincon - Constrained nonlinear minimization`
- Algorithm:** `Interior point`
- Problem:**
 - Objective function:** `@bigtopleft`
 - Derivatives:** `Gradient supplied`
 - Start point:** `[-1,-1,-1]`
- Constraints:**
 - Linear inequalities:** A: b:
 - Linear equalities:** Aeq: beq:
 - Bounds:** Lower: Upper:
 - Nonlinear constraint function:** `@twocone`
 - Derivatives:** `Gradient supplied`


- 2 For iterative output, scroll to the bottom of the **Options** pane and select **Level of display**, iterative.



3 In the **Options** pane, give the analytic Hessian function handle.



4 Under **Run solver and view results**, click **Start**.



To perform the minimization at the command line:

1 Set options as follows:

```

options =
optimoptions(@fmincon,'Algorithm','interior-point',...
'Display','off','GradObj','on','GradConstr','on',...
'Hessian','user-supplied','HessFcn',@hessinterior);

```

2 Run `fmincon` with starting point $[-1,-1,-1]$, using the options structure:

```

[x fval mflag output] = fmincon(@bigtopleft,[-1,-1,-1],...
[],[],[],[],[],[],@twocone,options)

```

The output is:

```

x =
-6.5000 -0.0000 -3.5000

```



```

fval =
    -2.8941e+03

mflag =
     1

output =
    iterations: 6
    funcCount: 7
    constrviolation: 0
    stepsize: 3.0479e-05
    algorithm: 'interior-point'
    firstorderopt: 5.9812e-05
    cgiterations: 3
    message: 'Local minimum found that satisfies the constraints.

Optimization complet...'

```

If you do not use a Hessian function, `fmincon` takes 9 iterations to converge, instead of 6:

```

options = optimoptions(@fmincon,'Algorithm','interior-point',...
    'Display','off','GradObj','on','GradConstr','on');

[x fval mflag output]=fmincon(@bigtopleft,[-1,-1,-1],...
    [],[],[],[],[],[],@twocone,options)

x =
    -6.5000    -0.0000    -3.5000

fval =
    -2.8941e+03

mflag =
     1

output =
    iterations: 9
    funcCount: 13

```

```
constrviolation: 2.9391e-08
  stepsize: 6.4842e-04
  algorithm: 'interior-point'
firstorderopt: 1.4235e-04
cgiterations: 0
  message: 'Local minimum found that satisfies the constraints.'
```

Optimization complet...'

Both runs lead to similar solutions, but the F-count and number of iterations are lower when using an analytic Hessian.

Nonlinear Equality and Inequality Constraints

For routines that permit equality constraints, nonlinear equality constraints must be computed with the nonlinear inequality constraints. For linear equalities, the coefficients of the equalities are passed in through the matrix Aeq and the right-hand-side vector beq.

For example, if you have the nonlinear equality constraint $x_1^2 + x_2 = 1$ and the nonlinear inequality constraint $x_1x_2 \geq -10$, rewrite them as

$$\begin{aligned}x_1^2 + x_2 - 1 &= 0, \\ -x_1x_2 - 10 &\leq 0,\end{aligned}$$

and then solve the problem using the following steps.

Step 1: Write a file objfun.m.

```
function f = objfun(x)
f = exp(x(1))*(4*x(1)^2+2*x(2)^2+4*x(1)*x(2)+2*x(2)+1);
```

Step 2: Write a file confuneq.m for the nonlinear constraints.

```
function [c, ceq] = confuneq(x)
% Nonlinear inequality constraints
c = -x(1)*x(2) - 10;
% Nonlinear equality constraints
ceq = x(1)^2 + x(2) - 1;
```

Step 3: Invoke constrained optimization routine.

```
x0 = [-1,1]; % Make a starting guess at the solution
options = optimoptions(@fmincon,'Algorithm','sqp');
[x,fval] = fmincon(@objfun,x0,[],[],[],[],[],[],...
    @confuneq,options);
```

After 21 function evaluations, the solution produced is

```
x, fval
x =
    -0.7529    0.4332
fval =
    1.5093

[c, ceq] = confuneq(x) % Check the constraint values at x

c =
    -9.6739
ceq =
    5.3291e-15
```

Note that `ceq` is equal to 0 within the default tolerance on the constraints of $1.0e-006$ and that `c` is less than or equal to 0, as desired.

Optimization App with the fmincon Solver

This example shows how to use the Optimization app with the `fmincon` solver to minimize a quadratic subject to linear and nonlinear constraints and bounds.

Consider the problem of finding $[x_1, x_2]$ that solves

$$\min_x f(x) = x_1^2 + x_2^2$$

subject to the constraints

$$\begin{array}{ll} 0.5 \leq x_1 & \text{(bound)} \\ -x_1 - x_2 + 1 \leq 0 & \text{(linear inequality)} \\ \left. \begin{array}{l} -x_1^2 - x_2^2 + 1 \leq 0 \\ -9x_1^2 - x_2^2 + 9 \leq 0 \\ -x_1^2 + x_2 \leq 0 \\ -x_2^2 + x_1 \leq 0 \end{array} \right\} & \text{(nonlinear inequality)} \end{array}$$

The starting guess for this problem is $x_1 = 3$ and $x_2 = 1$.

Step 1: Write a file `objecfun.m` for the objective function.

```
function f = objecfun(x)
f = x(1)^2 + x(2)^2;
```

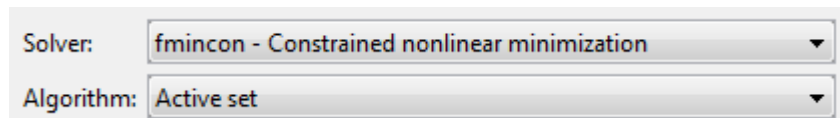
Step 2: Write a file `nonlconstr.m` for the nonlinear constraints.

```
function [c,ceq] = nonlconstr(x)
c = [-x(1)^2 - x(2)^2 + 1;
     -9*x(1)^2 - x(2)^2 + 9;
     -x(1)^2 + x(2);
     -x(2)^2 + x(1)];
```

```
ceq = [];
```

Step 3: Set up and run the problem with the Optimization app.

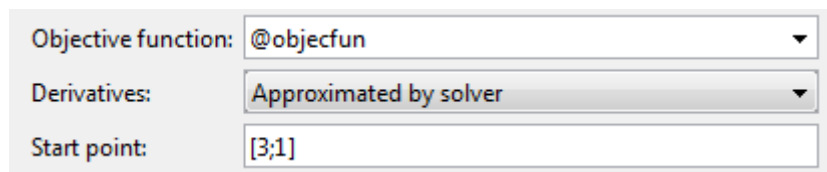
- 1 Enter `optimtool` in the Command Window to open the Optimization app.
- 2 Select `fmincon` from the selection of solvers and change the **Algorithm** field to `Active set`.



Solver: `fmincon - Constrained nonlinear minimization`

Algorithm: `Active set`

- 3 Enter `@objecfun` in the **Objective function** field to call the `objecfun.m` file.
- 4 Enter `[3;1]` in the **Start point** field.



Objective function: `@objecfun`

Derivatives: `Approximated by solver`

Start point: `[3;1]`

- 5 Define the constraints.
 - Set the bound $0.5 \leq x_1$ by entering `[0.5, -Inf]` in the **Lower** field. The `-Inf` entry means there is no lower bound on x_2 .
 - Set the linear inequality constraint by entering `[-1 -1]` in the **A** field and enter `-1` in the **b** field.
 - Set the nonlinear constraints by entering `@nonlconstr` in the **Nonlinear constraint function** field.

Constraints:

Linear inequalities: A: [-1 -1] b: -1

Linear equalities: Aeq: beq:

Bounds: Lower: [0.5, -Inf] Upper:

Nonlinear constraint function: @nonlconstr

Derivatives: Approximated by solver

- 6 In the **Options** pane, expand the **Display to command window** option if necessary, and select **Iterative** to show algorithm information at the Command Window for each iteration.

[-] Display to command window

Level of display: iterative

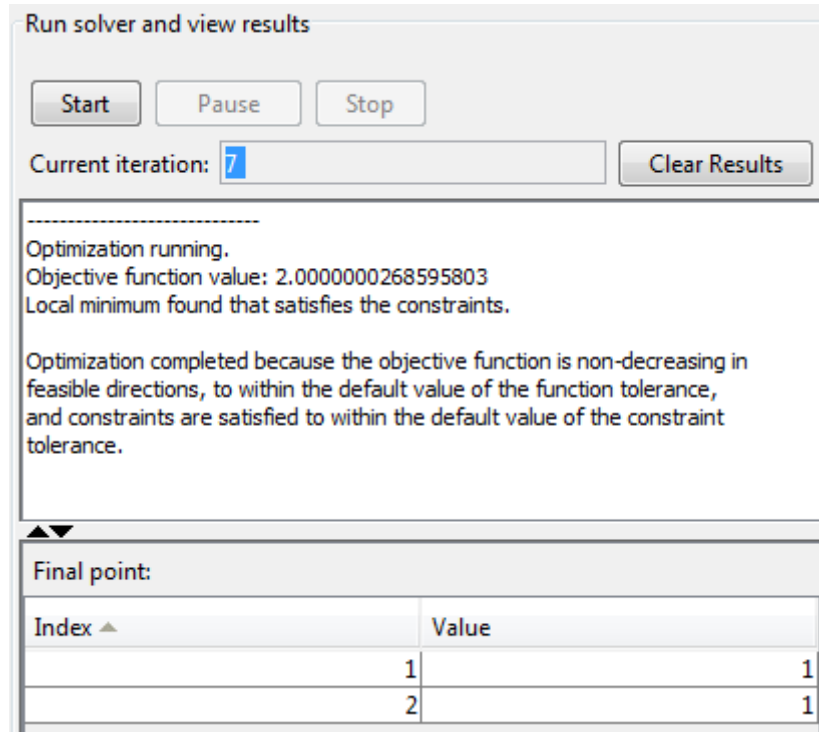
- 7 Click the **Start** button as shown in the following figure.

Run solver and view results

Start Pause Stop

Current iteration: Clear Results

- 8 When the algorithm terminates, under **Run solver and view results** the following information is displayed:



- The **Current iteration** value when the algorithm terminated, which for this example is 7.
- The final value of the objective function when the algorithm terminated:
Objective function value: 2.0000000268595803
- The algorithm termination message:

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.
- The final point, which for this example is

1

9 In the Command Window, the algorithm information is displayed for each iteration:

Iter	F-count	f(x)	Max constraint	Line search steplength	Directional derivative	First-order optimality	Procedure
0	3	10	2				Infeasible start point
1	6	4.84298	-0.1322	1	-5.22	1.74	
2	9	4.0251	-0.01168	1	-4.39	4.08	Hessian modified twice
3	12	2.42704	-0.03214	1	-3.85	1.09	
4	15	2.03615	-0.004728	1	-3.04	0.995	Hessian modified twice
5	18	2.00033	-5.596e-005	1	-2.82	0.0664	Hessian modified twice
6	21	2	-5.327e-009	1	-2.81	0.000522	Hessian modified twice

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.

Active inequalities (to within options.TolCon = 1e-006):

lower	upper	ineqlin	ineqnonlin
			3
			4

Minimization with Bound Constraints and Banded Preconditioner

The goal in this problem is to minimize the nonlinear function

$$f(x) = 1 + \sum_{i=1}^n |(3 - 2x_i)x_i - x_{i-1} - x_{i+1} + 1|^p + \sum_{i=1}^{n/2} |x_i + x_{i+n/2}|^p,$$

such that $-10.0 \leq x_i \leq 10.0$, where n is 800 (n should be a multiple of 4), $p = 7/3$, and $x_0 = x_{n+1} = 0$.


Step 1: Write a file `tbroyfg.m` that computes the objective function and the gradient of the objective

The `tbroyfg.m` file computes the function value and gradient. This file is long and is not included here. You can see the code for this function using the command

```
type tbroyfg
```


The sparsity pattern of the Hessian matrix has been predetermined and stored in the file `tbroyhstr.mat`. The sparsity structure for the Hessian of this problem is banded, as you can see in the following spy plot.

```
load tbroyhstr
spy(Hstr)
```



In this plot, the center stripe is itself a five-banded matrix. The following plot shows the matrix more clearly:

```
spy(Hstr(1:20,1:20))
```



Use `optimoptions` to set the `HessPattern` parameter to `Hstr`. When a problem as large as this has obvious sparsity structure, not setting the `HessPattern` parameter requires a huge amount of unnecessary memory and computation. This is because `fmincon` attempts to use finite differencing on a full Hessian matrix of 640,000 nonzero entries.

You must also set the `GradObj` parameter to `'on'` using `optimoptions`, since the gradient is computed in `tbroyfg.m`. Then execute `fmincon` as shown in Step 2.

Step 2: Call a nonlinear minimization routine with a starting point `xstart`.

```
fun = @tbroyfg;
load tbroyhstr          % Get Hstr, structure of the Hessian
n = 800;
xstart = -ones(n,1); xstart(2:2:n) = 1;
lb = -10*ones(n,1); ub = -lb;
options = optimoptions('fmincon','GradObj','on','HessPattern',Hstr,...
    'Algorithm','trust-region-reflective');

[x,fval,exitflag,output] = ...
    fmincon(fun,xstart,[],[],[],[],lb,ub,[],options);
```

After seven iterations, the `exitflag`, `fval`, and `output` values are

```
exitflag =
    3
```

```
fval =
    270.4790
```

```
output =
    iterations: 7
    funcCount: 8
    cgiterations: 18
    firstorderopt: 0.0163
    algorithm: 'trust-region-reflective'
    message: 'Local minimum possible.'
```

```
fmincon stopped because the final change in func...'  
    constrviolation: 0
```

For bound constrained problems, the first-order optimality is the infinity norm of $v \cdot *g$, where v is defined as in “Box Constraints” on page 6-30, and g is the gradient.

Because of the five-banded center stripe, you can improve the solution by using a five-banded preconditioner instead of the default diagonal preconditioner. Using the `optimoptions` function, reset the `PrecondBandWidth` parameter to 2 and solve the problem again. (The bandwidth is the number of upper (or lower) diagonals, not counting the main diagonal.)

```
fun = @tbroyfg;  
load tbroyhstr           % Get Hstr, structure of the Hessian  
n = 800;  
xstart = -ones(n,1); xstart(2:2:n,1) = 1;  
lb = -10*ones(n,1); ub = -lb;  
options = optimoptions('fmincon','GradObj','on','HessPattern',Hstr, ...  
    'Algorithm','trust-region-reflective','PrecondBandWidth',2);  
[x,fval,exitflag,output] = ...  
    fmincon(fun,xstart,[],[],[],[],lb,ub,[],options);
```

The number of iterations actually goes up by two; however the total number of CG iterations drops from 18 to 15. The first-order optimality measure is reduced by a factor of $1e-3$:

```
exitflag =  
    3  
  
fval =  
    270.4790  
  
output =  
    iterations: 9  
    funcCount: 10  
    cgiterations: 15  
    firstorderopt: 7.5340e-05  
    algorithm: 'trust-region-reflective'  
    message: 'Local minimum possible.'
```

```
fmincon stopped because the final change in func...'  
constrviolation: 0
```

Minimization with Linear Equality Constraints

The trust-region reflective method for `fmincon` can handle linear equality constraints if no other constraints exist. Suppose you want to minimize

$$f(x) = \sum_{i=1}^{n-1} \left((x_i^2)^{(x_{i+1}^2+1)} + (x_{i+1}^2)^{(x_i^2+1)} \right)$$

subject to some linear equality constraints. The objective function is coded in the function `brownfgh.m`. This example takes $n = 1000$. Furthermore, the `browneq.mat` file contains matrices `Aeq` and `beq` that represent the linear constraints $Aeq x = beq$. `Aeq` has 100 rows representing 100 linear constraints (so `Aeq` is a 100-by-1000 matrix).

Step 1: Write a file `brownfgh.m` that computes the objective function, the gradient of the objective, and the sparse tridiagonal Hessian matrix.

The file is lengthy so is not included here. View the code with the command

```
type brownfgh
```

Because `brownfgh` computes the gradient and Hessian values as well as the objective function, you need to use `optimoptions` to indicate that this information is available in `brownfgh`, using the `GradObj` and `Hessian` options.

The sparse matrix `Aeq` and vector `beq` are available in the file `browneq.mat`:

```
load browneq
```

The linear constraint system is 100-by-1000, has unstructured sparsity (use `spy(Aeq)` to view the sparsity structure), and is not too badly ill-conditioned:

```
condtest(Aeq*Aeq')
ans =
    2.9310e+006
```

Step 2: Call a nonlinear minimization routine with a starting point `xstart`.

```
fun = @brownfgh;
load browneq          % Get Aeq and beq, the linear equalities
n = 1000;
xstart = -ones(n,1); xstart(2:2:n) = 1;
options = optimoptions('fmincon','GradObj','on','Hessian','user-supplied',.
    'Algorithm','trust-region-reflective');
[x,fval,exitflag,output] = ...
    fmincon(fun,xstart,[],[],Aeq,beq,[],[],[],options);
```

`fmincon` prints the following exit message:

```
Local minimum possible.
```

```
fmincon stopped because the final change in function value relative to
its initial value is less than the default value of the function tolerance.
```

The `exitflag` value of 3 also indicates that the algorithm terminated because the change in the objective function value was less than the tolerance `TolFun`. The final function value is given by `fval`.

```
exitflag,fval,output
```

```
exitflag =
    3
```

```
fval =
    205.9313
```

```
output =
    iterations: 22
    funcCount: 23
    cgiterations: 30
    firstorderopt: 0.0027
    algorithm: 'trust-region-reflective'
    message: 'Local minimum possible.'
```

```
fmincon stopped because the final change in func...
constrviolation: 2.2293e-13
```


The linear equalities are satisfied at x .

$\text{norm}(A_{\text{eq}}*x - \text{beq})$

ans =

1.1919e-12

Minimization with Dense Structured Hessian, Linear Equalities

In this section...

“Hessian Multiply Function for Lower Memory” on page 6-80

“Step 1: Write a file brownvv.m that computes the objective function, the gradient, and the sparse part of the Hessian.” on page 6-81

“Step 2: Write a function to compute Hessian-matrix products for H given a matrix Y .” on page 6-81

“Step 3: Call a nonlinear minimization routine with a starting point and linear equality constraints.” on page 6-82

“Preconditioning” on page 6-84

Hessian Multiply Function for Lower Memory

The `fmincon` interior-point and trust-region-reflective algorithms, and the `fminunc` trust-region algorithm can solve problems where the Hessian is dense but structured. For these problems, `fmincon` and `fminunc` do not compute H^*Y with the Hessian H directly, because forming H would be memory-intensive. Instead, you must provide `fmincon` or `fminunc` with a function that, given a matrix Y and information about H , computes $W = H^*Y$.

In this example, the objective function is nonlinear and linear equalities exist so `fmincon` is used. The description applies to the trust-region reflective algorithm; the `fminunc` trust-region algorithm is similar. For the interior-point algorithm, see the 'HessMult' option in “Hessian” on page 10-49. The objective function has the structure

$$f(x) = \hat{f}(x) - \frac{1}{2}x^T VV^T x,$$

where V is a 1000-by-2 matrix. The Hessian of f is dense, but the Hessian of \hat{f} is sparse. If the Hessian of \hat{f} is \hat{H} , then H , the Hessian of f , is

$$H = \hat{H} - VV^T.$$

To avoid excessive memory usage that could happen by working with H directly, the example provides a Hessian multiply function, `hmf1eq1`. This function, when passed a matrix Y , uses sparse matrices `Hinfo`, which corresponds to \hat{H} , and V to compute the Hessian matrix product

$$W = H*Y = (\text{Hinfo} - V*V')*Y$$

In this example, the Hessian multiply function needs \hat{H} and V to compute the Hessian matrix product. V is a constant, so you can capture V in a function handle to an anonymous function.

However, \hat{H} is not a constant and must be computed at the current x . You can do this by computing \hat{H} in the objective function and returning \hat{H} as `Hinfo` in the third output argument. By using `optimoptions` to set the 'Hessian' options to 'on', `fmincon` knows to get the `Hinfo` value from the objective function and pass it to the Hessian multiply function `hmf1eq1`.

Step 1: Write a file `brownvv.m` that computes the objective function, the gradient, and the sparse part of the Hessian.

The example passes `brownvv` to `fmincon` as the objective function. The `brownvv.m` file is long and is not included here. You can view the code with the command

```
type brownvv
```

Because `brownvv` computes the gradient and part of the Hessian as well as the objective function, the example (Step 3) uses `optimoptions` to set the `GradObj` and `Hessian` options to 'on'.

Step 2: Write a function to compute Hessian-matrix products for H given a matrix Y .

Now, define a function `hmf1eq1` that uses `Hinfo`, which is computed in `brownvv`, and V , which you can capture in a function handle to an anonymous function, to compute the Hessian matrix product W where $W = H*Y = (\text{Hinfo} - V*V')*Y$. This function must have the form

```
W = hmfleq1(Hinfo,Y)
```

The first argument must be the same as the third argument returned by the objective function `brownvv`. The second argument to the Hessian multiply function is the matrix `Y` (of $W = H*Y$).

Because `fmincon` expects the second argument `Y` to be used to form the Hessian matrix product, `Y` is always a matrix with `n` rows where `n` is the number of dimensions in the problem. The number of columns in `Y` can vary. Finally, you can use a function handle to an anonymous function to capture `V`, so `V` can be the third argument to `'hmfleq1'`.

```
function W = hmfleq1(Hinfo,Y,V);  
%HMFLEQ1 Hessian-matrix product function for BROWNVV objective.  
% W = hmfleq1(Hinfo,Y,V) computes W = (Hinfo-V*V')*Y  
% where Hinfo is a sparse matrix computed by BROWNVV  
% and V is a 2 column matrix.  
W = Hinfo*Y - V*(V'*Y);
```

Note The function `hmfleq1` is available in the `optimdemos` folder as `hmfleq1.m`.

Step 3: Call a nonlinear minimization routine with a starting point and linear equality constraints.

Load the problem parameter, `V`, and the sparse equality constraint matrices, `Aeq` and `beq`, from `fleq1.mat`, which is available in the `optimdemos` folder. Use `optimoptions` to set the `GradObj` and `Hessian` options to `'on'` and to set the `HessMult` option to a function handle that points to `hmfleq1`. Call `fmincon` with objective function `brownvv` and with `V` as an additional parameter:

```
function [fval, exitflag, output, x] = runfleq1  
% RUNFLEQ1 demonstrates 'HessMult' option for FMINCON with linear  
% equalities.  
  
problem = load('fleq1'); % Get V, Aeq, beq  
V = problem.V; Aeq = problem.Aeq; beq = problem.beq;  
n = 1000; % problem dimension  
xstart = -ones(n,1); xstart(2:2:n,1) = ones(length(2:2:n),1); % starting po
```

```

options = optimoptions(@fmincon,'Algorithm','trust-region-reflective','Grad
    'Hessian','user-supplied','HessMult',@(Hinfo,Y)hmfleq1(Hinfo,Y,V),'Disp
    'TolFun',1e-9);
[x,fval,exitflag,output] = fmincon(@(x)brownvv(x,V),xstart,[],[],Aeq,beq,[],
    [],options);

```

To run the preceding code, enter

```
[fval,exitflag,output,x] = runfleq1;
```

Because the iterative display was set using `optimoptions`, this command generates the following iterative display:

Iteration	f(x)	Norm of step	First-order optimality	CG-iterations
0	1997.07		916	
1	1072.57	6.31716	465	1
2	480.247	8.19711	201	2
3	136.982	10.3039	78.1	2
4	44.416	9.04685	16.7	2
5	44.416	100	16.7	2
6	44.416	25	16.7	0
7	-9.05631	6.25	52.9	0
8	-317.437	12.5	91.7	1
9	-405.381	12.5	1.11e+003	1
10	-451.161	3.125	327	4
11	-482.688	0.78125	303	5
12	-547.427	1.5625	187	5
13	-610.42	1.5625	251	7
14	-711.522	1.5625	143	3
15	-802.98	3.125	165	3
16	-820.431	1.13329	32.9	3
17	-822.996	0.492813	7.61	2
18	-823.236	0.223154	1.68	3
19	-823.245	0.056205	0.529	3
20	-823.246	0.0150139	0.0342	5
21	-823.246	0.00479085	0.0152	7
22	-823.246	0.00353697	0.00828	9
23	-823.246	0.000884242	0.005	9
24	-823.246	0.0012715	0.00125	9

```
25          -823.246    0.000317876    0.0025    9
```

Local minimum possible.

fmincon stopped because the final change in function value relative to its initial value is less than the selected value of the function tolerance.

Convergence is rapid for a problem of this size with the PCG iteration cost increasing modestly as the optimization progresses. Feasibility of the equality constraints is maintained at the solution.

```
problem = load('fleq1'); % Get V, Aeq, beq
V = problem.V; Aeq = problem.Aeq; beq = problem.beq;
norm(Aeq*x-beq,inf)

ans =
    2.4869e-14
```

Preconditioning

In this example, `fmincon` cannot use `H` to compute a preconditioner because `H` only exists implicitly. Instead of `H`, `fmincon` uses `Hinfo`, the third argument returned by `brownvv`, to compute a preconditioner. `Hinfo` is a good choice because it is the same size as `H` and approximates `H` to some degree. If `Hinfo` were not the same size as `H`, `fmincon` would compute a preconditioner based on some diagonal scaling matrices determined from the algorithm. Typically, this would not perform as well.

Symbolic Math Toolbox Calculates Gradients and Hessians

If you have a Symbolic Math Toolbox license, you can easily calculate analytic gradients and Hessians for objective and constraint functions. There are two relevant Symbolic Math Toolbox functions:

- `jacobian` generates the gradient of a scalar function, and generates a matrix of the partial derivatives of a vector function. So, for example, you can obtain the Hessian matrix, the second derivatives of the objective function, by applying `jacobian` to the gradient. Part of this example shows how to use `jacobian` to generate symbolic gradients and Hessians of objective and constraint functions.
- `matlabFunction` generates either an anonymous function or a file that calculates the values of a symbolic expression. This example shows how to use `matlabFunction` to generate files that evaluate the objective and constraint function and their derivatives at arbitrary points.


Consider the electrostatics problem of placing 10 electrons in a conducting body. The electrons will arrange themselves so as to minimize their total potential energy, subject to the constraint of lying inside the body. It is well known that all the electrons will be on the boundary of the body at a minimum. The electrons are indistinguishable, so there is no unique minimum for this problem (permuting the electrons in one solution gives another valid solution). This example was inspired by Dolan, Moré, and Munson [58].

This example is a conducting body defined by the following inequalities:

$$z \leq -|x| - |y| \tag{6-61}$$

$$x^2 + y^2 + (z+1)^2 \leq 1. \tag{6-62}$$

This body looks like a pyramid on a sphere.



There is a slight gap between the upper and lower surfaces of the figure. This is an artifact of the general plotting routine used to create the figure. This routine erases any rectangular patch on one surface that touches the other surface.

The syntax and structures of the two sets of toolbox functions differ. In particular, symbolic variables are real or complex scalars, but Optimization Toolbox functions pass vector arguments. So there are several steps to take to generate symbolically the objective function, constraints, and all their requisite derivatives, in a form suitable for the interior-point algorithm of `fmincon`:

- 1 “Create the Variables” on page 6-87
- 2 “Include the Linear Constraints” on page 6-88
- 3 “Create the Nonlinear Constraints, Their Gradients and Hessians” on page 6-90
- 4 “Create the Objective Function, Its Gradient and Hessian” on page 6-91
- 5 “Create the Objective Function File” on page 6-92
- 6 “Create the Constraint Function File” on page 6-92
- 7 “Generate the Hessian Files” on page 6-93
- 8 “Run the Optimization” on page 6-94
- 9 “Clear the Symbolic Variable Assumptions” on page 6-99

To see the efficiency in using gradients and Hessians, see “Compare to Optimization Without Gradients and Hessians” on page 6-97.

Create the Variables

Generate a symbolic vector x as a 30-by-1 vector composed of real symbolic variables x_{ij} , i between 1 and 10, and j between 1 and 3. These variables represent the three coordinates of electron i : x_{i1} corresponds to the x coordinate, x_{i2} corresponds to the y coordinate, and x_{i3} corresponds to the z coordinate.

```
x = cell(3, 10);
for i = 1:10
    for j = 1:3
        x{j,i} = sprintf('x%d%d',i,j);
    end
end
x = x(:); % now x is a 30-by-1 vector
x = sym(x, 'real');
```

The vector x is:

```
x
```

x =
x11
x12
x13
x21
x22
x23
x31
x32
x33
x41
x42
x43
x51
x52
x53
x61
x62
x63
x71
x72
x73
x81
x82
x83
x91
x92
x93
x101
x102
x103

Include the Linear Constraints

Write the linear constraints in Equation 6-61,

$$z \leq -|x| - |y|,$$

as a set of four linear inequalities for each electron:

$$\begin{aligned}x_{i3} - x_{i1} - x_{i2} &\leq 0 \\x_{i3} - x_{i1} + x_{i2} &\leq 0 \\x_{i3} + x_{i1} - x_{i2} &\leq 0 \\x_{i3} + x_{i1} + x_{i2} &\leq 0\end{aligned}$$

Therefore there are a total of 40 linear inequalities for this problem.

Write the inequalities in a structured way:

$$B = [1, 1, 1; -1, 1, 1; 1, -1, 1; -1, -1, 1];$$

```
A = zeros(40,30);
for i=1:10
    A(4*i-3:4*i,3*i-2:3*i) = B;
end
```

$$b = \text{zeros}(40,1);$$

You can see that $A \cdot x \leq b$ represents the inequalities:

```
A*x
ans =
    x11 + x12 + x13
    x12 - x11 + x13
    x11 - x12 + x13
    x13 - x12 - x11
    x21 + x22 + x23
    x22 - x21 + x23
    x21 - x22 + x23
    x23 - x22 - x21
    x31 + x32 + x33
    x32 - x31 + x33
    x31 - x32 + x33
    x33 - x32 - x31
    x41 + x42 + x43
    x42 - x41 + x43
    x41 - x42 + x43
    x43 - x42 - x41
    x51 + x52 + x53
    x52 - x51 + x53
```

```

x51 - x52 + x53
x53 - x52 - x51
x61 + x62 + x63
x62 - x61 + x63
x61 - x62 + x63
x63 - x62 - x61
x71 + x72 + x73
x72 - x71 + x73
x71 - x72 + x73
x73 - x72 - x71
x81 + x82 + x83
x82 - x81 + x83
x81 - x82 + x83
x83 - x82 - x81
x91 + x92 + x93
x92 - x91 + x93
x91 - x92 + x93
x93 - x92 - x91
x101 + x102 + x103
x102 - x101 + x103
x101 - x102 + x103
x103 - x102 - x101

```

Create the Nonlinear Constraints, Their Gradients and Hessians

The nonlinear constraints in Equation 6-62 ,

$$x^2 + y^2 + (z+1)^2 \leq 1,$$

are also structured. Generate the constraints, their gradients, and Hessians as follows:

```

c = sym(zeros(1,10));
i = 1:10;
c = (x(3*i-2).^2 + x(3*i-1).^2 + (x(3*i)+1).^2 - 1).';

gradc = jacobian(c,x).'; % .' performs transpose

hessc = cell(1, 10);

```

```

for i = 1:10
    hessc{i} = jacobian(gradc(:,i),x);
end

```

The constraint vector **c** is a row vector, and the gradient of **c(i)** is represented in the *i*th column of the matrix **gradc**. This is the correct form, as described in “Nonlinear Constraints” on page 2-37.

The Hessian matrices, **hessc{1}...hessc{10}**, are square and symmetric. It is better to store them in a cell array, as is done here, than in separate variables such as **hessc1, ..., hessc10**.

Use the **.'** syntax to transpose. The **'** syntax means conjugate transpose, which has different symbolic derivatives.

Create the Objective Function, Its Gradient and Hessian

The objective function, potential energy, is the sum of the inverses of the distances between each electron pair:

$$\text{energy} = \sum_{i < j} \frac{1}{|x_i - x_j|}$$

The distance is the square root of the sum of the squares of the differences in the components of the vectors.

Calculate the energy, its gradient, and its Hessian as follows:

```

energy = sym(0);
for i = 1:3:25
    for j = i+3:3:28
        dist = x(i:i+2) - x(j:j+2);
        energy = energy + 1/sqrt(dist.'*dist);
    end
end

gradenergy = jacobian(energy,x).';

```

```
hessenergy = jacobian(gradenergy,x);
```

Create the Objective Function File

The objective function should have two outputs, energy and gradenergy. Put both functions in one vector when calling `matlabFunction` to reduce the number of subexpressions that `matlabFunction` generates, and to return the gradient only when the calling function (`fmincon` in this case) requests both outputs. This example shows placing the resulting files in your current folder. Of course, you can place them anywhere you like, as long as the folder is on the MATLAB path.

```
currdir = [pwd filesep]; % You may need to use currdir = pwd
filename = [currdir,'demoenergy.m'];
matlabFunction(energy,gradenergy,'file',filename,'vars',{x});
```

This syntax causes `matlabFunction` to return energy as the first output, and gradenergy as the second. It also takes a single input vector `{x}` instead of a list of inputs `x11, ..., x103`.

The resulting file `demoenergy.m` contains, in part, the following lines or similar ones:

```
function [energy,gradenergy] = demoenergy(in1)
%DEMOENERGY
% [ENERGY,GRADENERGY] = DEMOENERGY(IN1)
...
x101 = in1(28,:);
...
energy = 1./t140.^(1./2) + ...;
if nargout > 1
    ...
    gradenergy = [(t174.*(t185 - 2.*x11))./2 - ...];
end
```

This function has the correct form for an objective function with a gradient; see “Writing Scalar Objective Functions” on page 2-19.

Create the Constraint Function File

Generate the nonlinear constraint function, and put it in the correct format.

```
filename = [currdir, 'democonstr.m'];
matlabFunction(c, [], gradc, [], 'file', filename, 'vars', {x}, ...
    'outputs', {'c', 'ceq', 'gradc', 'gradceq'});
```

The resulting file `democonstr.m` contains, in part, the following lines or similar ones:

```
function [c,ceq,gradc,gradceq] = democonstr(in1)
%DEMOCONSTR
%    [C,CEQ,GRADC,GRADCEQ] = DEMOCONSTR(IN1)
...
x101 = in1(28,:);
...
c = [t417.^2 + ...];
if nargout > 1
    ceq = [];
end
if nargout > 2
    gradc = [2.*x11,...];
end
if nargout > 3
    gradceq = [];
end
```

This function has the correct form for a constraint function with a gradient; see “Nonlinear Constraints” on page 2-37.

Generate the Hessian Files

To generate the Hessian of the Lagrangian for the problem, first generate files for the energy Hessian and for the constraint Hessians.

The Hessian of the objective function, `hessenergy`, is a very large symbolic expression, containing over 150,000 symbols, as evaluating `size(char(hessenergy))` shows. So it takes a substantial amount of time to run `matlabFunction(hessenergy)`.

To generate a file `hessenergy.m`, run the following two lines:

```
filename = [currdir, 'hessenergy.m'];
```

```
matlabFunction(hessenergy, 'file', filename, 'vars', {x});
```

In contrast, the Hessians of the constraint functions are small, and fast to compute:

```
for i = 1:10
    ii = num2str(i);
    thename = ['hessc', ii];
    filename = [currdir, thename, '.m'];
    matlabFunction(hessc{i}, 'file', filename, 'vars', {x});
end
```

After generating all the files for the objective and constraints, put them together with the appropriate Lagrange multipliers in a file `hessfinal.m` as follows:

```
function H = hessfinal(X, lambda)
%
% Call the function hessenergy to start
H = hessenergy(X);

% Add the Lagrange multipliers * the constraint Hessians
H = H + hessc1(X) * lambda.ineqnonlin(1);
H = H + hessc2(X) * lambda.ineqnonlin(2);
H = H + hessc3(X) * lambda.ineqnonlin(3);
H = H + hessc4(X) * lambda.ineqnonlin(4);
H = H + hessc5(X) * lambda.ineqnonlin(5);
H = H + hessc6(X) * lambda.ineqnonlin(6);
H = H + hessc7(X) * lambda.ineqnonlin(7);
H = H + hessc8(X) * lambda.ineqnonlin(8);
H = H + hessc9(X) * lambda.ineqnonlin(9);
H = H + hessc10(X) * lambda.ineqnonlin(10);
```

Run the Optimization

Start the optimization with the electrons distributed randomly on a sphere of radius 1/2 centered at [0,0,-1]:

```
rng('default'); % for reproducibility
Xinitial = randn(3,10); % columns are normal 3-D vectors
for j=1:10
    Xinitial(:,j) = Xinitial(:,j)/norm(Xinitial(:,j))/2;
```



```
    % this normalizes to a 1/2-sphere
end
Xinitial(3,:) = Xinitial(3,:) - 1; % center at [0,0,-1]
Xinitial = Xinitial(:); % Convert to a column vector

Set the options to use the interior-point algorithm, and to use gradients and
the Hessian:

options =
optimoptions(@fmincon, 'Algorithm', 'interior-point', 'GradObj', 'on', ...
    'GradConstr', 'on', 'Hessian', 'user-supplied', ...
    'HessFcn', @hessfinal, 'Display', 'final');
```

Call `fmincon`:

```
[xfinal fval exitflag output] = fmincon(@demoenergy, Xinitial, ...
    A, b, [], [], [], [], @democonstr, options)
```

The output is as follows:

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in
feasible directions, to within the default value of the function tolerance,
and constraints are satisfied to within the default value of the constraint tolerance.

```
xfinal =
-0.0317
 0.0317
-1.9990
 0.6356
-0.6356
-1.4381
 0.0000
-0.0000
-0.0000
 0.0000
-1.0000
-1.0000
 1.0000
-0.0000
```

```
-1.0000
-1.0000
-0.0000
-1.0000
 0.6689
 0.6644
-1.3333
-0.6667
 0.6667
-1.3333
 0.0000
 1.0000
-1.0000
-0.6644
-0.6689
-1.3333


fval =
    34.1365

exitflag =
     1

output =
    iterations: 19
    funcCount: 28
    constrviolation: 0
    stepsize: 4.0372e-005
    algorithm: 'interior-point'
    firstorderopt: 4.0015e-007
    cgiterations: 55
    message: 'Local minimum found that satisfies the constraints.'

Optimization complet...
```

Even though the initial positions of the electrons were random, the final positions are nearly symmetric:



Compare to Optimization Without Gradients and Hessians

The use of gradients and Hessians makes the optimization run faster and more accurately. To compare with the same optimization using no gradient or Hessian information, set the options not to use gradients and Hessians:

```
options = optimoptions(@fmincon,'Algorithm','interior-point',...  
    'Display','final');  
[xfinal2 fval2 exitflag2 output2] = fmincon(@demoenergy,Xinitial,...  
    A,b,[],[],[],[],@democonstr,options)
```

The output shows that `fmincon` found an equivalent minimum, but took more iterations and many more function evaluations to do so:

```
Local minimum found that satisfies the constraints.
```

```
Optimization completed because the objective function is non-decreasing in  
feasible directions, to within the default value of the function tolerance,  
and constraints are satisfied to within the default value of the constraint tolerance.
```

```
xfinal2 =  
    0.0000  
    1.0000  
   -1.0000  
    0.6689  
   -0.6644  
   -1.3334  
   -0.6644  
    0.6689  
   -1.3334  
    0.0000  
   -1.0000  
   -1.0000  
    0.6357  
    0.6357  
   -1.4380  
   -0.0317  
   -0.0317  
   -1.9990  
    1.0000  
    0.0000  
   -1.0000  
   -1.0000  
    0.0000  
   -1.0000  
    0.0000  
    0.0000  
   -0.0000  
   -0.6667  
   -0.6667  
   -1.3334
```

```
fval2 =  
    34.1365  
  
exitflag2 =  
    1  
  
output2 =  
    iterations: 78  
    funcCount: 2463  
    constrviolation: 0  
    stepsize: 3.5711e-07  
    algorithm: 'interior-point'  
    firstorderopt: 2.7591e-06  
    cgiterations: 0  
    message: 'Local minimum found that satisfies the constraints.'  
  
Optimization complet...'
```

In this run the number of function evaluations (in `output2.funcCount`) is 2432, compared to 28 (in `output.funcCount`) when using gradients and Hessian.

Clear the Symbolic Variable Assumptions

The symbolic variables in this example have the assumption, in the symbolic engine workspace, that they are real. To clear this assumption from the symbolic engine workspace, it is not sufficient to delete the variables. You must clear the variables using the syntax

```
syms x11 x12 x13 clear
```

or reset the symbolic engine using the command

```
reset(symengine)
```

After resetting the symbolic engine you should clear all symbolic variables from the MATLAB workspace with the `clear` command, or `clear variable_list`.

One-Dimensional Semi-Infinite Constraints

Find values of x that minimize

$$f(x) = (x_1 - 0.5)^2 + (x_2 - 0.5)^2 + (x_3 - 0.5)^2$$

where

$$K_1(x, w_1) = \sin(w_1 x_1) \cos(w_1 x_2) - \frac{1}{1000}(w_1 - 50)^2 - \sin(w_1 x_3) - x_3 \leq 1,$$

$$K_2(x, w_2) = \sin(w_2 x_2) \cos(w_2 x_1) - \frac{1}{1000}(w_2 - 50)^2 - \sin(w_2 x_3) - x_3 \leq 1,$$

for all values of w_1 and w_2 over the ranges

$$1 \leq w_1 \leq 100,$$

$$1 \leq w_2 \leq 100.$$

Note that the semi-infinite constraints are one-dimensional, that is, vectors. Because the constraints must be in the form $K_i(x, w_i) \leq 0$, you need to compute the constraints as

$$K_1(x, w_1) = \sin(w_1 x_1) \cos(w_1 x_2) - \frac{1}{1000}(w_1 - 50)^2 - \sin(w_1 x_3) - x_3 - 1 \leq 0,$$

$$K_2(x, w_2) = \sin(w_2 x_2) \cos(w_2 x_1) - \frac{1}{1000}(w_2 - 50)^2 - \sin(w_2 x_3) - x_3 - 1 \leq 0.$$

First, write a file that computes the objective function.

```
function f = myfun(x,s)
% Objective function
f = sum((x-0.5).^2);
```

Second, write a file `mycon.m` that computes the nonlinear equality and inequality constraints and the semi-infinite constraints.

```
function [c,ceq,K1,K2,s] = mycon(X,s)
% Initial sampling interval
if isnan(s(1,1)),
```

```

    s = [0.2 0; 0.2 0];
end
% Sample set
w1 = 1:s(1,1):100;
w2 = 1:s(2,1):100;

% Semi-infinite constraints
K1 = sin(w1*X(1)).*cos(w1*X(2)) - 1/1000*(w1-50).^2 - ...
    sin(w1*X(3))-X(3)-1;
K2 = sin(w2*X(2)).*cos(w2*X(1)) - 1/1000*(w2-50).^2 - ...
    sin(w2*X(3))-X(3)-1;

% No finite nonlinear constraints
c = []; ceq=[];

% Plot a graph of semi-infinite constraints
plot(w1,K1,'-',w2,K2,':')
title('Semi-infinite constraints')
drawnow

```

Then, invoke an optimization routine.

```

x0 = [0.5; 0.2; 0.3];    % Starting guess
[x,fval] = fseminf(@myfun,x0,2,@mycon);

```

After eight iterations, the solution is

```

x
x =
    0.6675
    0.3012
    0.4022

```

The function value and the maximum values of the semi-infinite constraints at the solution x are

```

fval
fval =
    0.0771

```


```

[c,ceq,K1,K2] = mycon(x,NaN); % Initial sampling interval

```

```
max(K1)
ans =
    -0.0077
max(K2)
ans =
    -0.0812
```

A plot of the semi-infinite constraints is produced.



This plot shows how peaks in both constraints are on the constraint boundary.

The plot command inside `mycon.m` slows down the computation. Remove this line to improve the speed.

Two-Dimensional Semi-Infinite Constraint

Find values of x that minimize

$$f(x) = (x_1 - 0.2)^2 + (x_2 - 0.2)^2 + (x_3 - 0.2)^2,$$

where

$$K_1(x, w) = \sin(w_1 x_1) \cos(w_2 x_2) - \frac{1}{1000}(w_1 - 50)^2 - \sin(w_1 x_3) - x_3 + \dots$$

$$\sin(w_2 x_2) \cos(w_1 x_1) - \frac{1}{1000}(w_2 - 50)^2 - \sin(w_2 x_3) - x_3 \leq 1.5,$$

for all values of w_1 and w_2 over the ranges

$$1 \leq w_1 \leq 100,$$

$$1 \leq w_2 \leq 100,$$

starting at the point $x = [0.25, 0.25, 0.25]$.

Note that the semi-infinite constraint is two-dimensional, that is, a matrix.

First, write a file that computes the objective function.

```
function f = myfun(x,s)
% Objective function
f = sum((x-0.2).^2);
```

Second, write a file for the constraints, called `mycon.m`. Include code to draw the surface plot of the semi-infinite constraint each time `mycon` is called. This enables you to see how the constraint changes as X is being minimized.

```
function [c,ceq,K1,s] = mycon(X,s)
% Initial sampling interval
if isnan(s(1,1)),
    s = [2 2];
end

% Sampling set
w1x = 1:s(1,1):100;
```

```
w1y = 1:s(1,2):100;
[wx,wy] = meshgrid(w1x,w1y);

% Semi-infinite constraint
K1 = sin(wx*X(1)).*cos(wx*X(2))-1/1000*(wx-50).^2 -...
     sin(wx*X(3))-X(3)+sin(wy*X(2)).*cos(wx*X(1))-...
     1/1000*(wy-50).^2-sin(wy*X(3))-X(3)-1.5;

% No finite nonlinear constraints
c = []; ceq=[];

% Mesh plot
m = surf(wx,wy,K1,'edgecolor','none','facecolor','interp');
camlight headlight
title('Semi-infinite constraint')
drawnow
```

Next, invoke an optimization routine.

```
x0 = [0.25, 0.25, 0.25]; % Starting guess
[x,fval] = fseminf(@myfun,x0,1,@mycon)
```

After nine iterations, the solution is

```
x
x =
    0.2522    0.1714    0.1936
```

and the function value at the solution is


```
fval
fval =
    0.0036
```

The goal was to minimize the objective $f(x)$ such that the semi-infinite constraint satisfied $K_1(x,w) \leq 1.5$. Evaluating `mycon` at the solution `x` and looking at the maximum element of the matrix `K1` shows the constraint is easily satisfied.

```
[c,ceq,K1] = mycon(x,[0.5,0.5]); % Sampling interval 0.5
max(max(K1))
```

```
ans =  
-0.0333
```

This call to `mycon` produces the following surf plot, which shows the semi-infinite constraint at x .



Linear Programming Algorithms

In this section...

“Linear Programming Definition” on page 6-106

“Interior-Point Linear Programming” on page 6-106

“Active-Set linprog Algorithm” on page 6-110

“linprog Simplex Algorithm” on page 6-114

Linear Programming Definition

Linear programming is the problem of finding a vector x that minimizes a linear function $f^T x$ subject to linear constraints:

$$\min_x f^T x$$

such that one or more of the following hold: $Ax \leq b$, $Aeqx = beq$, $l \leq x \leq u$.

Interior-Point Linear Programming

Introduction

The default interior-point method is based on LIPSOL ([52]), which is a variant of Mehrotra’s predictor-corrector algorithm ([47]), a primal-dual interior-point method.

Main Algorithm

The algorithm begins by applying a series of preprocessing steps (see “Preprocessing” on page 6-109). After preprocessing, the problem has the form

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x = b \\ 0 \leq x \leq u. \end{cases} \quad (6-63)$$

The upper bounds constraints are implicitly included in the constraint matrix A . With the addition of primal slack variables s , Equation 6-63 becomes

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x = b \\ x + s = u \\ x \geq 0, s \geq 0. \end{cases} \quad (6-64)$$

which is referred to as the *primal* problem: x consists of the primal variables and s consists of the primal slack variables. The *dual* problem is

$$\max b^T y - u^T w \text{ such that } \begin{cases} A^T \cdot y - w + z = f \\ z \geq 0, w \geq 0, \end{cases} \quad (6-65)$$

where y and w consist of the dual variables and z consists of the dual slacks. The optimality conditions for this linear program, i.e., the primal Equation 6-64 and the dual Equation 6-65, are

$$F(x, y, z, s, w) = \begin{pmatrix} A \cdot x - b \\ x + s - u \\ A^T \cdot y - w + z - f \\ x_i z_i \\ s_i w_i \end{pmatrix} = \mathbf{0}, \quad x \geq 0, z \geq 0, s \geq 0, w \geq 0, \quad (6-66)$$

where $x_i z_i$ and $s_i w_i$ denote component-wise multiplication.

The quadratic equations $x_i z_i = 0$ and $s_i w_i = 0$ are called the *complementarity* conditions for the linear program; the other (linear) equations are called the *feasibility* conditions. The quantity

$$x^T z + s^T w$$

is the *duality gap*, which measures the residual of the complementarity portion of F when $(x, z, s, w) \geq 0$.

The algorithm is a *primal-dual algorithm*, meaning that both the primal and the dual programs are solved simultaneously. It can be considered a Newton-like method, applied to the linear-quadratic system $F(x, y, z, s, w) = 0$ in Equation 6-66, while at the same time keeping the iterates x , z , w , and s

positive, thus the name interior-point method. (The iterates are in the strictly interior region represented by the inequality constraints in Equation 6-64.)

The algorithm is a variant of the predictor-corrector algorithm proposed by Mehrotra. Consider an iterate $v = [x; y; z; s; w]$, where $[x; z; s; w] > 0$. First compute the so-called *prediction* direction

$$\Delta v_p = -\left(F^T(v)\right)^{-1} F(v),$$

which is the Newton direction; then the so-called *corrector* direction

$$\Delta v_c = -\left(F^T(v)\right)^{-1} F(v + \Delta v_p) - \mu e,$$

where $\mu > 0$ is called the *centering* parameter and must be chosen carefully.

\hat{e} is a zero-one vector with the ones corresponding to the quadratic equations in $F(v)$, i.e., the perturbations are only applied to the complementarity conditions, which are all quadratic, but not to the feasibility conditions, which are all linear. The two directions are combined with a step length parameter $\alpha > 0$ and update v to obtain the new iterate v^+ :

$$v^+ = v + \alpha (\Delta v_p + \Delta v_c),$$

where the step length parameter α is chosen so that

$$v^+ = [x^+; y^+; z^+; s^+; w^+]$$

satisfies

$$[x^+; z^+; s^+; w^+] > 0.$$

In solving for the preceding predictor/corrector directions, the algorithm computes a (sparse) direct factorization on a modification of the Cholesky factors of $A A^T$. If A has dense columns, it instead uses the Sherman-Morrison formula. If that solution is not adequate (the residual is too large), it performs an LDL factorization of an augmented system form of the step equations to

find a solution. (See Example 4 — The Structure of D in the MATLAB 1d1 function reference page.)

The algorithm then loops until the iterates converge. The main stopping criteria is a standard one:

$$\frac{\|r_b\|}{\max(1, \|b\|)} + \frac{\|r_f\|}{\max(1, \|f\|)} + \frac{\|r_u\|}{\max(1, \|u\|)} + \frac{|f^T x - b^T y + u^T w|}{\max(1, |f^T x|, |b^T y - u^T w|)} \leq tol,$$

where

$$r_b = Ax - b$$

$$r_f = A^T y - w + z - f$$

$$r_u = x + s - u$$

are the primal residual, dual residual, and upper-bound feasibility respectively, and

$$f^T x - b^T y + u^T w$$

is the difference between the primal and dual objective values, and *tol* is some tolerance. The sum in the stopping criteria measures the total relative errors in the optimality conditions in Equation 6-66.

Preprocessing

A number of preprocessing steps occur before the actual iterative algorithm begins. The resulting transformed problem is one where

- All variables are bounded below by zero.
- All constraints are equalities.
- Fixed variables, those with equal upper and lower bounds, are removed.
- Rows of all zeros in the constraint matrix are removed.
- The constraint matrix has full structural rank.

- Columns of all zeros in the constraint matrix are removed.
- When a significant number of singleton rows exist in the constraint matrix, the associated variables are solved for and the rows removed.

While these preprocessing steps can do much to speed up the iterative part of the algorithm, if the Lagrange multipliers are required, the preprocessing steps must be undone since the multipliers calculated during the algorithm are for the transformed problem, and not the original. Thus, if the multipliers are *not* requested, this transformation back is not computed, and might save some time computationally.

Active-Set linprog Algorithm

The medium-scale active-set linear programming algorithm is a variant of the sequential quadratic programming algorithm used by `fmincon` (“Sequential Quadratic Programming (SQP)” on page 6-33). The difference is that the quadratic term is set to zero.

At each major iteration of the SQP method, a QP problem of the following form is solved, where A_i refers to the i th row of the m -by- n matrix A .

$$\begin{aligned} \min_{d \in \mathcal{R}^n} q(d) &= c^T d, \\ A_i d &= b_i, \quad i = 1, \dots, m_e \\ A_i d &\leq b_i, \quad i = m_e + 1, \dots, m. \end{aligned}$$

The method used in Optimization Toolbox functions is an active set strategy (also known as a projection method) similar to that of Gill et al., described in [18] and [17]. It has been modified for both Linear Programming (LP) and Quadratic Programming (QP) problems.

The solution procedure involves two phases. The first phase involves the calculation of a feasible point (if one exists). The second phase involves the generation of an iterative sequence of feasible points that converge to the solution. In this method an active set, \bar{A}_k , is maintained that is an estimate of the active constraints (i.e., those that are on the constraint boundaries) at the solution point. Virtually all QP algorithms are active set methods. This

point is emphasized because there exist many different methods that are very similar in structure but that are described in widely different terms.

\bar{A}_k is updated at each iteration k , and this is used to form a basis for a search direction \hat{d}_k . Equality constraints always remain in the active set \bar{A}_k . The notation for the variable \hat{d}_k is used here to distinguish it from d_k in the major iterations of the SQP method. The search direction \hat{d}_k is calculated and minimizes the objective function while remaining on any active constraint boundaries. The feasible subspace for \hat{d}_k is formed from a basis Z_k whose columns are orthogonal to the estimate of the active set \bar{A}_k (i.e., $\bar{A}_k Z_k = 0$). Thus a search direction, which is formed from a linear summation of any combination of the columns of Z_k , is guaranteed to remain on the boundaries of the active constraints.

The matrix Z_k is formed from the last $m - l$ columns of the QR decomposition of the matrix \bar{A}_k^T , where l is the number of active constraints and $l < m$. That is, Z_k is given by

$$Z_k = Q[:, l+1 : m], \quad (6-67)$$

where

$$Q^T \bar{A}_k^T = \begin{bmatrix} R \\ 0 \end{bmatrix}.$$

Once Z_k is found, a new search direction \hat{d}_k is sought that minimizes $q(d)$ where \hat{d}_k is in the null space of the active constraints. That is, \hat{d}_k is a linear combination of the columns of Z_k : $\hat{d}_k = Z_k p$ for some vector p .

Then if you view the quadratic as a function of p , by substituting for \hat{d}_k , you have

$$q(p) = \frac{1}{2} p^T Z_k^T H Z_k p + c^T Z_k p. \quad (6-68)$$

Differentiating this with respect to p yields

$$\nabla q(p) = Z_k^T H Z_k p + Z_k^T c. \quad (6-69)$$

$\nabla q(p)$ is referred to as the projected gradient of the quadratic function because it is the gradient projected in the subspace defined by Z_k . The term $Z_k^T H Z_k$ is called the projected Hessian. Assuming the Hessian matrix H is positive definite (which is the case in this implementation of SQP), then the minimum of the function $q(p)$ in the subspace defined by Z_k occurs when $\nabla q(p) = 0$, which is the solution of the system of linear equations

$$Z_k^T H Z_k p = -Z_k^T c. \quad (6-70)$$

A step is then taken of the form

$$x_{k+1} = x_k + \alpha \hat{d}_k, \quad \text{where } \hat{d}_k = Z_k^T p. \quad (6-71)$$

At each iteration, because of the quadratic nature of the objective function, there are only two choices of step length α . A step of unity along \hat{d}_k is the exact step to the minimum of the function restricted to the null space of \bar{A}_k . If such a step can be taken, without violation of the constraints, then this is the solution to QP (Equation 6-68). Otherwise, the step along \hat{d}_k to the nearest constraint is less than unity and a new constraint is included in the active set at the next iteration. The distance to the constraint boundaries in any direction \hat{d}_k is given by

$$\alpha = \min_{i \in \{1, \dots, m\}} \left\{ \frac{-(A_i x_k - b_i)}{A_i \hat{d}_k} \right\}, \quad (6-72)$$

which is defined for constraints not in the active set, and where the direction \hat{d}_k is towards the constraint boundary, i.e., $A_i \hat{d}_k > 0$, $i = 1, \dots, m$.

When n independent constraints are included in the active set, without location of the minimum, Lagrange multipliers, λ_k , are calculated that satisfy the nonsingular set of linear equations

$$\bar{A}_k^T \lambda_k = c. \quad (6-73)$$

If all elements of λ_k are positive, x_k is the optimal solution of QP (Equation 6-68). However, if any component of λ_k is negative, and the component does not correspond to an equality constraint, then the corresponding element is deleted from the active set and a new iterate is sought.

Initialization

The algorithm requires a feasible point to start. If the current point from the SQP method is not feasible, then you can find a point by solving the linear programming problem

$$\begin{aligned} \min_{\gamma \in \mathcal{R}, x \in \mathcal{R}^n} \quad & \gamma \text{ such that} \\ A_i x = b_i, \quad & i = 1, \dots, m_e \\ A_i x - \gamma \leq b_i, \quad & i = m_e + 1, \dots, m. \end{aligned} \quad (6-74)$$

The notation A_i indicates the i th row of the matrix A . You can find a feasible point (if one exists) to Equation 6-74 by setting x to a value that satisfies the equality constraints. You can determine this value by solving an under- or overdetermined set of linear equations formed from the set of equality constraints. If there is a solution to this problem, then the slack variable γ is set to the maximum inequality constraint at this point.

You can modify the preceding QP algorithm for LP problems by setting the search direction to the steepest descent direction at each iteration, where g_k is the gradient of the objective function (equal to the coefficients of the linear objective function).

$$\hat{d}_k = -Z_k Z_k^T g_k. \quad (6-75)$$

If a feasible point is found using the preceding LP method, the main QP phase is entered. The search direction \hat{d}_k is initialized with a search direction \hat{d}_1 found from solving the set of linear equations

$$H\hat{d}_1 = -g_k, \quad (6-76)$$

where g_k is the gradient of the objective function at the current iterate x_k (i.e., $Hx_k + c$).

If a feasible solution is not found for the QP problem, the direction of search for the main SQP routine \hat{d}_k is taken as one that minimizes γ .

linprog Simplex Algorithm

The simplex algorithm, invented by George Dantzig in 1947, is one of the earliest and best known optimization algorithms. The algorithm solves the linear programming problem

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ lb \leq x \leq ub. \end{cases}$$

The algorithm moves along the edges of the polyhedron defined by the constraints, from one vertex to another, while decreasing the value of the objective function, $f^T x$, at each step. This section describes an improved version of the original simplex algorithm that returns a vertex optimal solution.

This section covers the following topics:

- “Main Algorithm” on page 6-114
- “Preprocessing” on page 6-116
- “Using the Simplex Algorithm” on page 6-117
- “Basic and Nonbasic Variables” on page 6-117

Main Algorithm

The simplex algorithm has two phases:

- Phase 1 — Compute an initial basic feasible point.

- Phase 2 — Compute the optimal solution to the original problem.

Note You cannot supply an initial point x_0 for `linprog` with the simplex algorithm. If you pass in x_0 as an input argument, `linprog` ignores x_0 and computes its own initial point for the algorithm.

Phase 1. In phase 1, the algorithm finds an initial basic feasible solution (see “Basic and Nonbasic Variables” on page 6-117 for a definition) by solving an auxiliary piecewise linear programming problem. The objective function of

the auxiliary problem is the *linear penalty function* $P = \sum_j P_j(x_j)$,

where $P_j(x_j)$ is defined by

$$P_j(x_j) = \begin{cases} x_j - u_j & \text{if } x_j > u_j \\ 0 & \text{if } l_j \leq x_j \leq u_j \\ l_j - x_j & \text{if } l_j > x_j. \end{cases}$$

$P(x)$ measures how much a point x violates the lower and upper bound conditions. The auxiliary problem is

$$\min_x \sum_j P_j \quad \text{subject to} \quad \begin{cases} A \cdot x \leq b \\ A_{eq} \cdot x = b_{eq}. \end{cases}$$

The original problem has a feasible basis point if and only if the auxiliary problem has minimum value 0.

The algorithm finds an initial point for the auxiliary problem by a heuristic method that adds slack and artificial variables as necessary. The algorithm then uses this initial point together with the simplex algorithm to solve the auxiliary problem. The solution is the initial point for phase 2 of the main algorithm.

Phase 2. In phase 2, the algorithm applies the simplex algorithm, starting at the initial point from phase 1, to solve the original problem. At each iteration, the algorithm tests the optimality condition and stops if the current solution is optimal. If the current solution is not optimal, the algorithm

- 1 Chooses one variable, called the *entering variable*, from the nonbasic variables and adds the corresponding column of the nonbasis to the basis (see “Basic and Nonbasic Variables” on page 6-117 for definitions).
- 2 Chooses a variable, called the *leaving variable*, from the basic variables and removes the corresponding column from the basis.
- 3 Updates the current solution and the current objective value.

The algorithm chooses the entering and the leaving variables by solving two linear systems while maintaining the feasibility of the solution.

The algorithm detects when there is no progress in the Phase 2 solution process. It attempts to continue by performing bound perturbation. For an explanation of this part of the algorithm, see Applegate, Bixby, Chvatal, and Cook [59].

Preprocessing

The simplex algorithm uses the same preprocessing steps as the interior-point linear programming solver, which are described in “Preprocessing” on page 6-109. In addition, the algorithm uses two other steps:

- 1 Eliminates columns that have only one nonzero element and eliminates their corresponding rows.
- 2 For each constraint equation $a x = b$, where a is a row of A_{eq} , the algorithm computes the lower and upper bounds of the linear combination $a x$ as rlb and rub if the lower and upper bounds are finite. If either rlb or rub equals b , the constraint is called a *forcing constraint*. The algorithm sets each variable corresponding to a nonzero coefficient of $a x$ equal to its upper or lower bound, depending on the forcing constraint. The algorithm then deletes the columns corresponding to these variables and deletes the rows corresponding to the forcing constraints.

Using the Simplex Algorithm

To use the simplex method, set the Algorithm option to 'simplex' using `optimoptions`.

```
options = optimoptions(@linprog,'Algorithm','simplex');
```

Then call `linprog` with the options input argument. See the reference page for `linprog` for more information.

`linprog` returns empty output arguments for `x` and `fval` if it detects infeasibility or unboundedness in the preprocessing procedure. `linprog` returns the current point when it

- Exceeds the maximum number of iterations
- Detects that the problem is infeasible or unbounded in phases 1 or 2

When the problem is unbounded, `linprog` returns `x` and `fval` in the unbounded direction.

Basic and Nonbasic Variables

This section defines the terms *basis*, *nonbasis*, and *basic feasible solutions* for a linear programming problem. The definition assumes that the problem is given in the following standard form:

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x = b, \\ lb \leq x \leq ub. \end{cases}$$

(Note that A and b are not the matrix and vector defining the inequalities in the original problem.) Assume that A is an m -by- n matrix, of rank $m < n$,

whose columns are $\{a_1, a_2, \dots, a_n\}$. Suppose that $\{a_{i_1}, a_{i_2}, \dots, a_{i_m}\}$ is a basis for the column space of A , with index set $B = \{i_1, i_2, \dots, i_m\}$, and that $N = \{1, 2, \dots, n\} \setminus B$ is the complement of B . The submatrix A_B is called a *basis* and the complementary submatrix A_N is called a *nonbasis*. The vector of *basic variables* is x_B and the vector of *nonbasic variables* is x_N . At each iteration in phase 2, the algorithm replaces one column of the current basis with a column of the nonbasis and updates the variables x_B and x_N accordingly.

If x is a solution to $Ax = b$ and all the nonbasic variables in x_N are equal to either their lower or upper bounds, x is called a *basic solution*. If, in addition, the basic variables in x_B satisfy their lower and upper bounds, so that x is a feasible point, x is called a *basic feasible solution*.

Typical Linear Programming Problem

This example shows the solution of a typical linear programming problem. The problem is

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ x \geq 0. \end{cases}$$

You can load the matrices and vectors A, Aeq, b, beq, f, and the lower bounds lb into the MATLAB workspace with

```
load sc50b
```

This problem in sc50b.mat has 48 variables, 30 inequalities, and 20 equalities.

Use linprog to solve the problem:

```
options = optimoptions(@linprog,'Display','iter');
[x,fval,exitflag,output] = ...
    linprog(f,A,b,Aeq,beq,lb,[],[],options);
```

Because the iterative display was set using optimoptions, the results displayed are

Residuals:	Primal	Dual	Duality	Total
	Infeas	Infeas	Gap	Rel
	A*x-b	A'*y+z-f	x'*z	Error

Iter 0:	1.50e+03	2.19e+01	1.91e+04	1.00e+02
Iter 1:	1.15e+02	3.16e-15	3.62e+03	9.90e-01
Iter 2:	9.79e-13	2.62e-15	4.32e+02	9.48e-01
Iter 3:	3.49e-12	5.93e-15	7.78e+01	6.88e-01
Iter 4:	4.86e-11	8.35e-16	2.38e+01	2.69e-01
Iter 5:	2.18e-10	3.39e-16	5.05e+00	6.89e-02
Iter 6:	1.05e-10	9.55e-17	1.64e-01	2.34e-03
Iter 7:	9.43e-12	1.51e-16	1.09e-05	1.55e-07
Iter 8:	1.11e-12	1.68e-16	1.09e-11	1.52e-13

Optimization terminated.

For this problem, the interior-point linear programming algorithm quickly reduces the scaled residuals below the default tolerance of $1e-08$.

The `exitflag` value is positive, telling you `linprog` converged. You can also get the final function value in `fval` and the number of iterations in `output.iterations`:

```
exitflag,fval,output
```

```
exitflag =  
    1
```

```
fval =  
-70.0000
```

```
output =  
    iterations: 8  
    algorithm: 'interior-point'  
    cgiterations: 0  
    message: 'Optimization terminated.'  
    constrviolation: 4.8317e-13  
    firstorderopt: 2.7908e-13
```

Quadratic Programming Algorithms

In this section...

“Quadratic Programming Definition” on page 6-121

“interior-point-convex quadprog Algorithm” on page 6-121

“trust-region-reflective quadprog Algorithm” on page 6-126

“active-set quadprog Algorithm” on page 6-131

Quadratic Programming Definition

Quadratic programming is the problem of finding a vector x that minimizes a quadratic function, possibly subject to linear constraints:

$$\min_x \frac{1}{2} x^T H x + c^T x$$

such that $A x \leq b$, $A_{eq} x = b_{eq}$, $l \leq x \leq u$.

interior-point-convex quadprog Algorithm

The interior-point-convex algorithm performs the following steps:

- 1 “Presolve/Postsolve” on page 6-121
- 2 “Generate Initial Point” on page 6-122
- 3 “Predictor-Corrector” on page 6-123
- 4 “Multiple Corrections” on page 6-125
- 5 “Total Relative Error” on page 6-125

Presolve/Postsolve

The algorithm begins by attempting to simplify the problem by removing redundancies and simplifying constraints. In particular, the presolve portion performs the following tasks, among others:

- Check if any variables have equal upper and lower bounds. If so, check for feasibility, then fix and remove the variables.
- Check if any linear inequality constraints involve just one variable. If so, check for feasibility, and change the linear constraint to a bound.
- Check if any linear equality constraints involve just one variable. If so, check for feasibility, then fix and remove the variable.
- Check if any linear constraint matrix has zero rows. If so, check for feasibility, and delete the rows.
- Check if the bounds and linear constraints are consistent.
- Check if any variables appear only as linear terms in the objective function and do not appear in any linear constraint. If so, check for feasibility and boundedness, and fix the variables at their appropriate bounds.

In the presolve step the algorithm might detect an infeasible or unbounded problem. If so, the algorithm halts and issues an appropriate exit message.

If the algorithm does not detect an infeasible or unbounded problem in the presolve step, it continues with the other steps. At the end, it reconstructs the original problem, undoing any presolve transformations. This final step is the postsolve step.

For details, see Gould and Toint [63].

Generate Initial Point

The initial point x_0 for the algorithm is:

- 1** Initialize x_0 to $\text{ones}(n, 1)$, where n is the number of rows in H .
- 2** For components that have both an upper bound ub and a lower bound lb , if a component of x_0 is not strictly inside the bounds, the component is set to $(ub + lb) / 2$.
- 3** For components that have only one bound, modify the component if necessary to lie strictly inside the bound.

Predictor-Corrector

Similar to the `fmincon` interior-point algorithm, the interior-point-convex algorithm tries to find a point where the Karush-Kuhn-Tucker (KKT) conditions hold. For the quadratic programming problem described in “Quadratic Programming Definition” on page 6-121, these conditions are:

$$\begin{aligned} Hx + c - A_{eq}^T y - \bar{A}^T z &= 0 \\ \bar{A}x - \bar{b} - s &= 0 \\ A_{eq}x - b_{eq} &= 0 \\ s_i z_i &= 0, \quad i = 1, 2, \dots, m \\ s &\geq 0 \\ z &\geq 0. \end{aligned}$$

Here

- \bar{A} is the extended linear inequality matrix that includes bounds written as linear inequalities. \bar{b} is the corresponding linear inequality vector, including bounds.
- s is the vector of slacks that convert inequality constraints to equalities. s has length m , the number of linear inequalities and bounds.
- z is the vector of Lagrange multipliers corresponding to s .
- y is the vector of Lagrange multipliers associated with the equality constraints.

The algorithm first predicts a step from the Newton-Raphson formula, then computes a corrector step. The corrector attempts to better enforce the nonlinear constraint $s_i z_i = 0$.

Definitions for the predictor step:

- r_d , the dual residual:

$$r_d = Hx + c - A_{eq}^T y - \bar{A}^T z.$$

- r_{eq} , the primal equality constraint residual:

$$r_{eq} = A_{eq}x - b_{eq}.$$

- r_{ineq} , the primal inequality constraint residual, which includes bounds and slacks:

$$r_{ineq} = \bar{A}x - \bar{b} - s.$$

- r_{sz} , the complementarity residual:

$$r_{sz} = Sz.$$

S is the diagonal matrix of slack terms, z is the column matrix of Lagrange multipliers.

- r_c , the average complementarity:

$$r_c = \frac{s^T z}{m}.$$

In a Newton step, the changes in x , s , y , and z , are given by:

$$\begin{pmatrix} H & 0 & -A_{eq}^T & -\bar{A}^T \\ A_{eq} & 0 & 0 & 0 \\ \bar{A} & -I & 0 & 0 \\ 0 & Z & 0 & S \end{pmatrix} \begin{pmatrix} \Delta x \\ \Delta s \\ \Delta y \\ \Delta z \end{pmatrix} = - \begin{pmatrix} r_d \\ r_{eq} \\ r_{ineq} \\ r_{sz} \end{pmatrix}.$$

However, a full Newton step might be infeasible, because of the positivity constraints on s and z . Therefore, quadprog shortens the step, if necessary, to maintain positivity.

Additionally, to maintain a “centered” position in the interior, instead of trying to solve $s_i z_i = 0$, the algorithm takes a positive parameter σ , and tries to solve

$$s_i z_i = \sigma r_c.$$

quadprog replaces r_{sz} in the Newton step equation with $r_{sz} + \Delta s \Delta z - \sigma r_c \mathbf{1}$, where $\mathbf{1}$ is the vector of ones. Also, quadprog reorders the Newton equations

to obtain a symmetric, more numerically stable system for the predictor step calculation.

For details, see Mehrotra [47].

Multiple Corrections

After calculating the corrected Newton step, quadprog can perform more calculations to get both a longer current step, and to prepare for better subsequent steps. These multiple correction calculations can improve both performance and robustness. For details, see Gondzio [62].

Total Relative Error

quadprog calculates a *merit function* φ at every iteration. The merit function is a measure of feasibility, and is also called *total relative error*. quadprog stops if the merit function grows too large. In this case, quadprog declares the problem to be infeasible.

The merit function is related to the KKT conditions for the problem—see “Predictor-Corrector” on page 6-123. Use the following definitions:

$$\begin{aligned}\rho &= \max\left(1, \|H\|, \|\bar{A}\|, \|A_{\text{eq}}\|, \|c\|, \|\bar{b}\|, \|b_{\text{eq}}\|\right) \\ r_{\text{eq}} &= A_{\text{eq}}x - b_{\text{eq}} \\ r_{\text{ineq}} &= \bar{A}x - \bar{b} + s \\ r_{\text{d}} &= Hx + c + A_{\text{eq}}^T \lambda_{\text{eq}} + \bar{A}^T \bar{\lambda}_{\text{ineq}} \\ g &= x^T Hx + f^T x - \bar{b}^T \bar{\lambda}_{\text{ineq}} - b_{\text{eq}}^T \lambda_{\text{eq}}.\end{aligned}$$

The notation \bar{A} and \bar{b} means the linear inequality coefficients, augmented with terms to represent bounds. The notation $\bar{\lambda}_{\text{ineq}}$ similarly represents Lagrange multipliers for the linear inequality constraints, including bound constraints. This was called z in “Predictor-Corrector” on page 6-123, and

λ_{eq} was called y .

The merit function φ is

$$\frac{1}{\rho} \left(\max \left(\|r_{\text{eq}}\|_{\infty}, \|r_{\text{ineq}}\|_{\infty}, \|r_{\text{d}}\|_{\infty} \right) + g \right).$$

quadprog iterative display includes a column showing the merit function under the heading Total relative error.

trust-region-reflective quadprog Algorithm

Many of the methods used in Optimization Toolbox solvers are based on *trust regions*, a simple yet powerful concept in optimization.

To understand the trust-region approach to optimization, consider the unconstrained minimization problem, minimize $f(x)$, where the function takes vector arguments and returns scalars. Suppose you are at a point x in n -space and you want to improve, i.e., move to a point with a lower function value. The basic idea is to approximate f with a simpler function q , which reasonably reflects the behavior of function f in a neighborhood N around the point x . This neighborhood is the trust region. A trial step s is computed by minimizing (or approximately minimizing) over N . This is the trust-region subproblem,

$$\min_s \{q(s), s \in N\}. \tag{6-77}$$

The current point is updated to be $x + s$ if $f(x + s) < f(x)$; otherwise, the current point remains unchanged and N , the region of trust, is shrunk and the trial step computation is repeated.

The key questions in defining a specific trust-region approach to minimizing $f(x)$ are how to choose and compute the approximation q (defined at the current point x), how to choose and modify the trust region N , and how accurately to solve the trust-region subproblem. This section focuses on the unconstrained problem. Later sections discuss additional complications due to the presence of constraints on the variables.

In the standard trust-region method ([48]), the quadratic approximation q is defined by the first two terms of the Taylor approximation to F at x ; the neighborhood N is usually spherical or ellipsoidal in shape. Mathematically the trust-region subproblem is typically stated

$$\min \left\{ \frac{1}{2} s^T H s + s^T g \text{ such that } \|D s\| \leq \Delta \right\}, \quad (6-78)$$

where g is the gradient of f at the current point x , H is the Hessian matrix (the symmetric matrix of second derivatives), D is a diagonal scaling matrix, Δ is a positive scalar, and $\| \cdot \|$ is the 2-norm. Good algorithms exist for solving Equation 6-78 (see [48]); such algorithms typically involve the computation of a full eigensystem and a Newton process applied to the secular equation

$$\frac{1}{\Delta} - \frac{1}{\|s\|} = 0.$$

Such algorithms provide an accurate solution to Equation 6-78. However, they require time proportional to several factorizations of H . Therefore, for large-scale problems a different approach is needed. Several approximation and heuristic strategies, based on Equation 6-78, have been proposed in the literature ([42] and [50]). The approximation approach followed in Optimization Toolbox solvers is to restrict the trust-region subproblem to a two-dimensional subspace S ([39] and [42]). Once the subspace S has been computed, the work to solve Equation 6-78 is trivial even if full eigenvalue/eigenvector information is needed (since in the subspace, the problem is only two-dimensional). The dominant work has now shifted to the determination of the subspace.

The two-dimensional subspace S is determined with the aid of a preconditioned conjugate gradient process described below. The solver defines S as the linear space spanned by s_1 and s_2 , where s_1 is in the direction of the gradient g , and s_2 is either an approximate Newton direction, i.e., a solution to

$$H \cdot s_2 = -g, \quad (6-79)$$

or a direction of negative curvature,

$$s_2^T \cdot H \cdot s_2 < 0. \quad (6-80)$$

The philosophy behind this choice of S is to force global convergence (via the steepest descent direction or negative curvature direction) and achieve fast local convergence (via the Newton step, when it exists).

A sketch of unconstrained minimization using trust-region ideas is now easy to give:

- 1 Formulate the two-dimensional trust-region subproblem.
- 2 Solve Equation 6-78 to determine the trial step s .
- 3 If $f(x + s) < f(x)$, then $x = x + s$.
- 4 Adjust Δ .

These four steps are repeated until convergence. The trust-region dimension Δ is adjusted according to standard rules. In particular, it is decreased if the trial step is not accepted, i.e., $f(x + s) \geq f(x)$. See [46] and [49] for a discussion of this aspect.

Optimization Toolbox solvers treat a few important special cases of f with specialized functions: nonlinear least-squares, quadratic functions, and linear least-squares. However, the underlying algorithmic ideas are the same as for the general case. These special cases are discussed in later sections.

The subspace trust-region method is used to determine a search direction. However, instead of restricting the step to (possibly) one reflection step, as in the nonlinear minimization case, a piecewise reflective line search is conducted at each iteration. See [45] for details of the line search.

Preconditioned Conjugate Gradient Method

A popular way to solve large symmetric positive definite systems of linear equations $Hp = -g$ is the method of Preconditioned Conjugate Gradients (PCG). This iterative approach requires the ability to calculate matrix-vector products of the form Hv where v is an arbitrary vector. The symmetric positive definite matrix M is a *preconditioner* for H . That is, $M = C^2$, where $C^{-1}HC^{-1}$ is a well-conditioned matrix or a matrix with clustered eigenvalues.

In a minimization context, you can assume that the Hessian matrix H is symmetric. However, H is guaranteed to be positive definite only in the neighborhood of a strong minimizer. Algorithm PCG exits when a direction of negative (or zero) curvature is encountered, i.e., $d^T Hd \leq 0$. The PCG output direction, p , is either a direction of negative curvature or an approximate (*tol* controls how approximate) solution to the Newton system $Hp = -g$. In

either case p is used to help define the two-dimensional subspace used in the trust-region approach discussed in “Trust-Region Methods for Nonlinear Minimization” on page 6-5.

Linear Equality Constraints

Linear constraints complicate the situation described for unconstrained minimization. However, the underlying ideas described previously can be carried through in a clean and efficient way. The trust-region methods in Optimization Toolbox solvers generate strictly feasible iterates.

The general linear equality constrained minimization problem can be written

$$\min\{f(x) \text{ such that } Ax = b\}, \quad (6-81)$$

where A is an m -by- n matrix ($m \leq n$). Some Optimization Toolbox solvers preprocess A to remove strict linear dependencies using a technique based on the LU factorization of A^T [46]. Here A is assumed to be of rank m .

The method used to solve Equation 6-81 differs from the unconstrained approach in two significant ways. First, an initial feasible point x_0 is computed, using a sparse least-squares step, so that $Ax_0 = b$. Second, Algorithm PCG is replaced with Reduced Preconditioned Conjugate Gradients (RPCG), see [46], in order to compute an approximate reduced Newton step (or a direction of negative curvature in the null space of A). The key linear algebra step involves solving systems of the form

$$\begin{bmatrix} C & \tilde{A}^T \\ \tilde{A} & 0 \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix}, \quad (6-82)$$

where \tilde{A} approximates A (small nonzeros of A are set to zero provided rank is not lost) and C is a sparse symmetric positive-definite approximation to H , i.e., $C = H$. See [46] for more details.

Box Constraints

The box constrained problem is of the form

$$\min\{f(x) \text{ such that } l \leq x \leq u\}, \quad (6-83)$$

where l is a vector of lower bounds, and u is a vector of upper bounds. Some (or all) of the components of l can be equal to $-\infty$ and some (or all) of the components of u can be equal to ∞ . The method generates a sequence of strictly feasible points. Two techniques are used to maintain feasibility while achieving robust convergence behavior. First, a scaled modified Newton step replaces the unconstrained Newton step (to define the two-dimensional subspace S). Second, reflections are used to increase the step size.

The scaled modified Newton step arises from examining the Kuhn-Tucker necessary conditions for Equation 6-83,

$$(D(x))^{-2} g = 0, \quad (6-84)$$

where

$$D(x) = \text{diag}\left(|v_k|^{-1/2}\right),$$

and the vector $v(x)$ is defined below, for each $1 \leq i \leq n$:

- If $g_i < 0$ and $u_i < \infty$ then $v_i = x_i - u_i$
- If $g_i \geq 0$ and $l_i > -\infty$ then $v_i = x_i - l_i$
- If $g_i < 0$ and $u_i = \infty$ then $v_i = -1$
- If $g_i \geq 0$ and $l_i = -\infty$ then $v_i = 1$

The nonlinear system Equation 6-84 is not differentiable everywhere. Nondifferentiability occurs when $v_i = 0$. You can avoid such points by maintaining strict feasibility, i.e., restricting $l < x < u$.

The scaled modified Newton step s_k for the nonlinear system of equations given by Equation 6-84 is defined as the solution to the linear system

$$\hat{M}D s^N = -\hat{g} \quad (6-85)$$

at the k th iteration, where

$$\hat{g} = D^{-1}g = \text{diag}(|v|^{1/2})g, \quad (6-86)$$

and

$$\hat{M} = D^{-1}HD^{-1} + \text{diag}(g)J^v. \quad (6-87)$$

Here J^v plays the role of the Jacobian of $|v|$. Each diagonal component of the diagonal matrix J^v equals 0, -1 , or 1 . If all the components of l and u are finite, $J^v = \text{diag}(\text{sign}(g))$. At a point where $g_i = 0$, v_i might not be differentiable.

$J_{ii}^v = 0$ is defined at such a point. Nondifferentiability of this type is not a cause for concern because, for such a component, it is not significant which value v_i takes. Further, $|v_i|$ will still be discontinuous at this point, but the function $|v_i|g_i$ is continuous.

Second, reflections are used to increase the step size. A (single) reflection step is defined as follows. Given a step p that intersects a bound constraint, consider the first bound constraint crossed by p ; assume it is the i th bound constraint (either the i th upper or i th lower bound). Then the reflection step $p^R = p$ except in the i th component, where $p_i^R = -p_i$.

active-set quadprog Algorithm

Recall the problem quadprog addresses:

$$\min_x \frac{1}{2}x^T Hx + c^T x \quad (6-88)$$

such that $Ax \leq b$, $Aeqx = beq$, and $l \leq x \leq u$. m is the total number of linear constraints, the sum of number of rows of A and of Aeq .

The medium-scale quadprog algorithm is an active-set strategy (also known as a projection method) similar to that of Gill et al., described in [18] and [17]. It has been modified for both Linear Programming (LP) and Quadratic Programming (QP) problems.

The solution procedure involves two phases. The first phase involves the calculation of a feasible point (if one exists). The second phase involves the

generation of an iterative sequence of feasible points that converge to the solution.

Active Set Iterations

In this method an active set matrix, S_k , is maintained that is an estimate of the active constraints (i.e., those that are on the constraint boundaries) at the solution point. Specifically, the active set S_k consists of the rows of A_{eq} , and a subset of the rows of A . S_k is updated at each iteration k , and is used to form a basis for a search direction d_k . Equality constraints always remain in the active set S_k . The search direction d_k is calculated and minimizes the objective function while remaining on active constraint boundaries. The feasible subspace for d_k is formed from a basis Z_k whose columns are orthogonal to the estimate of the active set S_k (i.e., $S_k Z_k = 0$). Thus a search direction, which is formed from a linear summation of any combination of the columns of Z_k , is guaranteed to remain on the boundaries of the active constraints.

The matrix Z_k is formed from the last $m - l$ columns of the QR decomposition of the matrix S_k^T , where l is the number of active constraints and $l < m$. That is, Z_k is given by

$$Z_k = Q[:, l+1 : m], \quad (6-89)$$

where

$$Q^T S_k^T = \begin{bmatrix} R \\ 0 \end{bmatrix}.$$

Once Z_k is found, a search direction d_k is sought that minimizes the objective function at d_k , where d_k is in the null space of the active constraints. That is, d_k is a linear combination of the columns of Z_k : $d_k = Z_k p$ for some vector p .

Then if you view the quadratic objective function as a function of p , by substituting for d_k , the result is

$$q(p) = \frac{1}{2} p^T Z_k^T H Z_k p + c^T Z_k p. \quad (6-90)$$

Differentiating this with respect to p yields

$$\nabla q(p) = Z_k^T H Z_k p + Z_k^T c. \quad (6-91)$$

$\nabla q(p)$ is referred to as the projected gradient of the quadratic function because it is the gradient projected in the subspace defined by Z_k . The term $Z_k^T H Z_k$ is called the projected Hessian. Assuming the Hessian matrix H is positive definite, the minimum of the function $q(p)$ in the subspace defined by Z_k occurs when $\nabla q(p) = 0$, which is the solution of the system of linear equations

$$Z_k^T H Z_k p = -Z_k^T c. \quad (6-92)$$

The next step is

$$x_{k+1} = x_k + \alpha d_k, \quad \text{where } d_k = Z_k^T p. \quad (6-93)$$

At each iteration, because of the quadratic nature of the objective function, there are only two choices of step length α . A step of unity along d_k is the exact step to the minimum of the function restricted to the null space of S_k . If such a step can be taken, without violation of the constraints, then this is the solution to QP (Equation 6-88). Otherwise, the step along d_k to the nearest constraint is less than unity and a new constraint is included in the active set at the next iteration. The distance to the constraint boundaries in any direction d_k is given by

$$\alpha = \min_{i \in \{1, \dots, m\}} \left\{ \frac{-(A_i x_k - b_i)}{A_i d_k} \right\}, \quad (6-94)$$

which is defined for constraints not in the active set, and where the direction d_k is towards the constraint boundary, i.e., $A_i d_k > 0$, $i = 1, \dots, m$.

Lagrange multipliers, λ_k , are calculated that satisfy the nonsingular set of linear equations

$$S_k^T \lambda_k = c. \quad (6-95)$$

If all elements of λ_k are positive, x_k is the optimal solution of QP (Equation 6-88). However, if any component of λ_k is negative, and the component does

not correspond to an equality constraint, then the corresponding element is deleted from the active set and a new iterate is sought.

Initialization

The algorithm requires a feasible point to start. If the initial point is not feasible, then you can find a feasible point by solving the linear programming problem

$$\begin{aligned} \min_{\gamma \in \mathfrak{R}, x \in \mathfrak{R}^n} \quad & \gamma \text{ such that} \\ A_i x = b_i, \quad & i = 1, \dots, m_e \text{ (the rows of } Aeq) \\ A_i x - \gamma \leq b_i, \quad & i = m_e + 1, \dots, m \text{ (the rows of } A). \end{aligned} \tag{6-96}$$

The notation A_i indicates the i th row of the matrix A . You can find a feasible point (if one exists) to Equation 6-96 by setting x to a value that satisfies the equality constraints. You can determine this value by solving an under- or overdetermined set of linear equations formed from the set of equality constraints. If there is a solution to this problem, the slack variable γ is set to the maximum inequality constraint at this point.

You can modify the preceding QP algorithm for LP problems by setting the search direction d to the steepest descent direction at each iteration, where g_k is the gradient of the objective function (equal to the coefficients of the linear objective function):

$$d = -Z_k Z_k^T g_k. \tag{6-97}$$

If a feasible point is found using the preceding LP method, the main QP phase is entered. The search direction d_k is initialized with a search direction d_1 found from solving the set of linear equations

$$Hd_1 = -g_k, \tag{6-98}$$

where g_k is the gradient of the objective function at the current iterate x_k (i.e., $Hx_k + c$).

Quadratic Minimization with Bound Constraints

To minimize a large-scale quadratic with upper and lower bounds, you can use the `quadprog` function with the 'trust-region-reflective' algorithm.

The problem stored in the MAT-file `qbox1.mat` is a positive definite quadratic, and the Hessian matrix `H` is tridiagonal, subject to upper (`ub`) and lower (`lb`) bounds.

Step 1: Load the Hessian and define `f`, `lb`, and `ub`.

```
load qbox1    % Get H
lb = zeros(400,1); lb(400) = -inf;
ub = 0.9*ones(400,1); ub(400) = inf;
f = zeros(400,1); f([1 400]) = -2;
```

Step 2: Call a quadratic minimization routine with a starting point `xstart`.

```
xstart = 0.5*ones(400,1);
options = optimoptions('quadprog','Algorithm','trust-region-reflective');
[x,fval,exitflag,output] = ...
    quadprog(H,f,[],[],[],[],lb,ub,xstart,options);
```

Looking at the resulting values of `exitflag` and `output`,

```
exitflag,output
```

```
exitflag =
    3
```

```
output =
```

```
    algorithm: 'trust-region-reflective'
    iterations: 19
   constrviolation: 0
   firstorderopt: 8.3903e-06
    cgiterations: 1673
    message: 'Optimization terminated: relative function value change
    less than s...'
```

You can see that while convergence occurred in 20 iterations, the high number of CG iterations indicates that the cost of the linear system solve is high. In light of this cost, one strategy would be to limit the number of CG iterations per optimization iteration. The default number is the dimension of the problem divided by two, 200 for this problem. Suppose you limit it to 50 using the `MaxPCGIter` flag in options:

```
options = optimoptions(options,'MaxPCGIter',50);
[x,fval,exitflag,output] = ...
    quadprog(H,f,[],[],[],[],lb,ub,xstart,options);
```

This time convergence still occurs and the total number of CG iterations (1547) has dropped:

```
exitflag,output
```

```
exitflag =
    3
```

```
output =
    algorithm: 'trust-region-reflective'
    iterations: 36
    constrviolation: 0
    firstorderopt: 2.3821e-005
    cgiterations: 1547
    message: 'Optimization terminated: relative function value change less than s...'
```

A second strategy would be to use a direct solver at each iteration by setting the `PrecondBandWidth` option to `inf`:

```
options = optimoptions(options,'PrecondBandWidth',inf);
[x,fval,exitflag,output] = ...
    quadprog(H,f,[],[],[],[],lb,ub,xstart,options);
```

Now the number of iterations has dropped to 10:

```
exitflag,output
```

```
exitflag =
    3
```

```

output =
    algorithm: 'trust-region-reflective'
    iterations: 10
    constrviolation: 0
    firstorderopt: 2.8219e-06
    cgiterations: 0
    message: 'Optimization terminated: relative function value chan
than s...'

```

Using a direct solver at each iteration usually causes the number of iterations to decrease, but often takes more time per iteration. For this problem, the tradeoff is beneficial, as the time for `quadprog` to solve the problem decreases by a factor of 10.

You can also use the default `'interior-point-convex'` algorithm to solve this convex problem:

```

options =
optimoptions('quadprog','Algorithm','interior-point-convex');
[x,fval,exitflag,output] = ...
    quadprog(H,f,[],[],[],[],lb,ub,[],options);

```

Check the exit flag and output structure:

```
exitflag,output
```

```
exitflag =
    1
```

```

output =
    message: 'Minimum found that satisfies the constraints.
Optimization completed bec...'
    algorithm: 'interior-point-convex'
    firstorderopt: 1.4120e-06
    constrviolation: 0
    iterations: 8
    cgiterations: []

```

Quadratic Minimization with Dense, Structured Hessian

The quadprog trust-region-reflective method can also solve large problems where the Hessian is dense but structured. For these problems, quadprog does not compute H^*Y with the Hessian H directly, as it does for active-set problems and for trust-region-reflective problems with sparse H , because forming H would be memory-intensive. Instead, you must provide quadprog with a function that, given a matrix Y and information about H , computes $W = H^*Y$.

In this example, the Hessian matrix H has the structure $H = B + A^*A'$ where B is a sparse 512-by-512 symmetric matrix, and A is a 512-by-10 sparse matrix composed of a number of dense columns. To avoid excessive memory usage that could happen by working with H directly because H is dense, the example provides a Hessian multiply function, `qpbox4mult`. This function, when passed a matrix Y , uses sparse matrices A and B to compute the Hessian matrix product $W = H^*Y = (B + A^*A')^*Y$.

In this example, the matrices A and B need to be provided to the Hessian multiply function `qpbox4mult`. You can pass one matrix as the first argument to `quadprog`, which is passed to the Hessian multiply function. You can use a nested function to provide the value of the second matrix.

Step 1: Decide what part of H to pass to `quadprog` as the first argument.

Either A or B can be passed as the first argument to `quadprog`. The example chooses to pass B as the first argument because this results in a better preconditioner (see “Preconditioning” on page 6-141).

```
quadprog(B, f, [], [], [], [], 1, u, xstart, options)
```

Step 2: Write a function to compute Hessian-matrix products for H .

Now, define a function `runqpbox4` that

- Contains a nested function `qpbox4mult` that uses A and B to compute the Hessian matrix product W , where $W = H^*Y = (B + A^*A')^*Y$. The nested function must have the form

```
W = qpbox4mult(Hinfo,Y,...)
```

The first two arguments `Hinfo` and `Y` are required.

- Loads the problem parameters from `qpbox4.mat`.
- Uses `optimoptions` to set the `HessMult` option to a function handle that points to `qpbox4mult`.
- Calls `quadprog` with `B` as the first argument.

The first argument to the nested function `qpbox4mult` must be the same as the first argument passed to `quadprog`, which in this case is the matrix `B`.

The second argument to `qpbox4mult` is the matrix `Y` (of $W = H*Y$). Because `quadprog` expects `Y` to be used to form the Hessian matrix product, `Y` is always a matrix with `n` rows, where `n` is the number of dimensions in the problem. The number of columns in `Y` can vary. The function `qpbox4mult` is nested so that the value of the matrix `A` comes from the outer function. Optimization Toolbox software includes the `runqpbox4.m` file.

```
function [fval, exitflag, output, x] = runqpbox4
%RUNQPBOX4 demonstrates 'HessMult' option for QUADPROG with bounds.

problem = load('qpbox4'); % Get xstart, u, l, B, A, f
xstart = problem.xstart; u = problem.u; l = problem.l;
B = problem.B; A = problem.A; f = problem.f;
mtxmpy = @qpbox4mult; % function handle to qpbox4mult nested function

% Choose algorithm and the HessMult option
options = optimoptions(@quadprog,'Algorithm','trust-region-reflective','HessMult',mtxmpy);

% Pass B to qpbox4mult via the H argument. Also, B will be used in
% computing a preconditioner for PCG.
[x, fval, exitflag, output] = quadprog(B,f,[],[],[],[],l,u,xstart,options);

function W = qpbox4mult(B,Y)
%QPBOX4MULT Hessian matrix product with dense structured Hessian.
% W = qpbox4mult(B,Y) computes W = (B + A*A')*Y where
% INPUT:
% B - sparse square matrix (512 by 512)
% Y - vector (or matrix) to be multiplied by B + A'*A.
```

```
% VARIABLES from outer function runqbbox4:
%     A - sparse matrix with 512 rows and 10 columns.
%
% OUTPUT:
%     W - The product (B + A*A')*Y.
%
% Order multiplies to avoid forming A*A',
%   which is large and dense
W = B*Y + A*(A'*Y);
end
```

```
end
```

Step 3: Call a quadratic minimization routine with a starting point.

To call the quadratic minimizing routine contained in runqbbox4, enter

```
[fval,exitflag,output] = runqbbox4;
```

to run the preceding code. Then display the values for fval, exitflag, and output. The results are

```
Optimization terminated: relative function value changing by
less than sqrt(OPTIONS.TolFun), no negative curvature detected
in current trust region model and the rate of progress (change
in f(x)) is slow.
```

```
fval,exitflag,output
```

```
fval =
-1.0538e+003
```

```
exitflag =
3
```

```
output =
    algorithm: 'trust-region-reflective'
    iterations: 18
    constrviolation: 0
```

```

    firstorderopt: 0.0043
    cgiterations: 30
    message: 'Optimization terminated: relative function value change
    less than s...'

```

After 18 iterations with a total of 30 PCG iterations, the function value is reduced to

```

fval
fval =
    -1.0538e+003

```

and the first-order optimality is

```

output.firstorderopt
ans =
    0.0043

```

Preconditioning

In this example, `quadprog` cannot use `H` to compute a preconditioner because `H` only exists implicitly. Instead, `quadprog` uses `B`, the argument passed in instead of `H`, to compute a preconditioner. `B` is a good choice because it is the same size as `H` and approximates `H` to some degree. If `B` were not the same size as `H`, `quadprog` would compute a preconditioner based on some diagonal scaling matrices determined from the algorithm. Typically, this would not perform as well.

Because the preconditioner is more approximate than when `H` is available explicitly, adjusting the `TolPcg` parameter to a somewhat smaller value might be required. This example is the same as the previous one, but reduces `TolPcg` from the default 0.1 to 0.01.

```

function [fval, exitflag, output, x] = runqpbox4prec
%RUNQPBOX4PREC demonstrates 'HessMult' option for QUADPROG with bounds.

problem = load('qpbox4'); % Get xstart, u, l, B, A, f
xstart = problem.xstart; u = problem.u; l = problem.l;
B = problem.B; A = problem.A; f = problem.f;
mtxmpy = @qpbox4mult; % function handle to qpbox4mult nested function

% Choose algorithm, the HessMult option, and override the TolPCG option

```

```

options = optimoptions(@quadprog,'Algorithm','trust-region-reflective',...
    'HessMult','mtxmpy','TolPcg',0.01);

% Pass B to qpbox4mult via the H argument. Also, B will be used in
% computing a preconditioner for PCG.
% A is passed as an additional argument after 'options'
[x, fval, exitflag, output] = quadprog(B,f,[],[],[],[],l,u,xstart,options);

function W = qpbox4mult(B,Y)
    %QPBOX4MULT Hessian matrix product with dense structured Hessian.
    % W = qpbox4mult(B,Y) computes  $W = (B + A^*A^*)^*Y$  where
    % INPUT:
    %     B - sparse square matrix (512 by 512)
    %     Y - vector (or matrix) to be multiplied by  $B + A^*A^*$ .
    % VARIABLES from outer function runqpbox4prec:
    %     A - sparse matrix with 512 rows and 10 columns.
    %
    % OUTPUT:
    %     W - The product  $(B + A^*A^*)^*Y$ .
    %
    % Order multiplies to avoid forming  $A^*A^*$ ,
    % which is large and dense
    W = B*Y + A*(A'^*Y);
end

end

```

Now, enter

```
[fval,exitflag,output] = runqpbox4prec;
```

to run the preceding code. After 18 iterations and 50 PCG iterations, the function value has the same value to five significant digits

```
fval
fval =
-1.0538e+003
```

but the first-order optimality is further reduced.


```
output.firstorderopt  
ans =  
    0.0028
```

Note Decreasing TolPcg too much can substantially increase the number of PCG iterations.

Large Sparse Quadratic Program with Interior Point Algorithm

This example shows the value of using sparse arithmetic when you have a sparse problem. The matrix has n rows, where you choose n . For some large n , the active-set algorithm runs out of memory, but the interior-point-convex algorithm works fine.

The problem is to minimize $x' * H * x / 2 + f' * x$ subject to

$$x(1) + x(2) + \dots + x(n) = 0,$$

$$\text{where } f = [-1; -2; -3; \dots; -n].$$

- 1 Create the parameter n and the utility matrix T . The matrix T is a sparse circulant matrix that is simply a helper for creating the sparse positive-definite quadratic matrix H .


```
n = 30000; % Adjust n to a large value
T = spalloc(n,n,n); % make a sparse circulant matrix
r = 1:n-1;
for m = r
    T(m,m+1)=1;
end
T(n,1) = 1;
```

- 2 Create a sparse vector v . Then create the matrix H by shifted versions of $v * v'$. The matrix T creates shifts of v .

```
v(n) = 0; v(1) = 1; v(2) = 2; v(4) = 3;
v = (sparse(v))';
% Make a banded type of matrix
H = spalloc(n,n,7*n);
r = 1:n;
for m = r
    H = H + v*v';
    v = T*v;
end
```

- 3 Take a look at the structure of H :

spy(H)



- 4** Create the problem vector f and linear constraint.

```
f = -r; % linear term
A = ones(1,n); b = 0;
```

- 5** Solve the quadratic programming problem with the interior-point-convex algorithm.

```
options =
optimoptions(@quadprog, 'Algorithm', 'interior-point-convex');
[x,fval,exitflag,output,lambda] = ...
quadprog(H,f,A,b,[],[],[],[],[],options);
```

Minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the selected

value of the function tolerance, and constraints are satisfied to within the selected value of the constraint tolerance.

- 6 View the solution value, output structure, and Lagrange multiplier:

```
fval,output,lambda
```

```
fval =  
-3.1331e+10
```

```
output =  
    message: 'Minimum found that satisfies the constraints.'
```

```
Optimization completed bec...'  
    algorithm: 'interior-point-convex'  
    firstorderopt: 1.2753e-04  
    constrviolation: 8.5020e-09  
    iterations: 6  
    cgiterations: []
```

```
lambda =  
    ineqlin: 1.5000e+004  
    eqlin: [0x1 double]  
    lower: [30000x1 double]  
    upper: [30000x1 double]
```

Since there are no lower bounds or upper bounds, all the values in `lambda.lower` and `lambda.upper` are 0. The inequality constraint is active, since `lambda.ineqlin` is nonzero.

- 7 Notice that `quadprog` with the active-set algorithm fails with an out-of-memory error:

```
options = optimoptions(@quadprog,'Algorithm','active-set');  
[x fval] = quadprog(H,f,A,b,[],[],[],[],[],options);
```

```
Warning: Cannot use sparse matrices with active-set algorithm: converting t  
> In quadprog at 409  
Error using full  
Out of memory. Type HELP MEMORY for your options.
```

```
Error in quadprog (line 410)  
    H = full(H); A = full(A); Aeq = full(Aeq);
```

Binary Integer Programming Algorithms

In this section...

“Binary Integer Programming Definition” on page 6-148

“bintprog Algorithm” on page 6-148

Binary Integer Programming Definition

Binary integer programming is the problem of finding a binary vector x that minimizes a linear function $f^T x$ subject to linear constraints:

$$\min_x f^T x$$

such that $Ax \leq b$, $Aeq\ x = beq$, x binary.

bintprog Algorithm

bintprog uses a linear programming (LP)-based branch-and-bound algorithm to solve binary integer programming problems. The algorithm searches for an optimal solution to the binary integer programming problem by solving a series of *LP-relaxation* problems, in which the binary integer requirement on the variables is replaced by the weaker constraint $0 \leq x \leq 1$. The algorithm


- Searches for a binary integer feasible solution
- Updates the best binary integer feasible point found so far as the search tree grows
- Verifies that no better integer feasible solution is possible by solving a series of linear programming problems

The following sections describe the branch-and-bound method in greater detail.

Branching

The algorithm creates a search tree by repeatedly adding constraints to the problem, that is, "branching." At a branching step, the algorithm chooses a

variable x_j , whose current value is not an integer and adds the constraint $x_j = 0$ to form one branch and the constraint $x_j = 1$ to form the other branch. This process can be represented by a binary tree, in which the nodes represent the added constraints. The following picture illustrates a complete binary tree for a problem that has three variables, x_1 , x_2 , and x_3 . Note that, in general, the order of the variables going down the levels in the tree is not the usual order of their subscripts



Deciding Whether to Branch

At each node, the algorithm solves an LP-relaxation problem using the constraints at that node and decides whether to branch or to move to another node depending on the outcome. There are three possibilities:

- If the LP-relaxation problem at the current node is infeasible or its optimal value is greater than that of the best integer point, the algorithm removes the node from the tree, after which it does not search any branches below that node. The algorithm then moves to a new node according to the method you specify in `NodeSearchStrategy` option.
- If the algorithm finds a new feasible integer point with lower objective value than that of the best integer point, it updates the current best integer point and moves to the next node.
- If the LP-relaxation problem is optimal but not integer and the optimal objective value of the LP relaxation problem is less than the best integer point, the algorithm branches according to the method you specify in the `BranchStrategy` option.

See “Options” on page 10-6 for a description of the `NodeSearchStrategy` and `BranchStrategy` options.

Bounds

The solution to the LP-relaxation problem provides a lower bound for the binary integer programming problem. If the solution to the LP-relaxation problem is already a binary integer vector, it provides an upper bound for the binary integer programming problem.

As the search tree grows more nodes, the algorithm updates the lower and upper bounds on the objective function, using the bounds obtained in the bounding step. The bound on the objective value serves as the threshold to cut off unnecessary branches.

Limits for the Algorithm

The algorithm for `binprog` could potentially search all 2^n binary integer vectors, where n is the number of variables. As a complete search might take a very long time, you can limit the search using the following options

- `MaxNodes` — Maximum number of nodes the algorithm searches
- `MaxRLPIter` — Maximum number of iterations the LP-solver performs at any node
- `MaxTime` — Maximum amount of time in seconds the algorithm runs

See “Options” on page 10-6 for more information.

Optimal Investments Via Binary Integer Programming

In this section...

“Investments with Constraints” on page 6-151

“Problem Statement” on page 6-151

“bintprog Formulation” on page 6-153

“bintprog Solution” on page 6-155

Investments with Constraints

This example uses bintprog to solve an integer programming problem that is not obviously a binary integer programming problem. This is done by representing each nonbinary integer-valued variable as an appropriate sum of binary variables, and by using linear constraints carefully. While the example is not particularly realistic, it demonstrates a variety of techniques:

- How to formulate nonbinary integer programming problems
- How to formulate an objective and constraints
- How to use indicator variables (y_i in the example)

Problem Statement

There are five investment opportunities labeled 1, 2, 3, 4, and 5. The investments have the costs and payoffs listed in the following table.

Investment	Buy-In Cost	Cost/Unit	Payoff/Unit	Max # Units
1	\$25	\$5	\$15	5
2	\$35	\$7	\$25	4
3	\$28	\$6	\$17	5
4	\$20	\$4	\$13	7
5	\$40	\$8	\$18	3

- The maximum total investment is \$125.
- The problem is to maximize profit, which is payoff minus cost.

- The payoff is the sum of the units bought times the payoff/unit.
- The cost per investment is the buy-in cost plus the cost/unit times the number of units if you buy at least one unit; otherwise, it is 0.
- The cost is the sum of the costs per investment.

It is convenient to formulate this problem using the indicator variables y_i . Define these as $y_i = 1$ when corresponding quantity variable x_i is positive, and $y_i = 0$ when $x_i = 0$:

- $x_i = \#$ units purchased of investment i
- $y_i = 1$ if $x_i > 0$, $y_i = 0$ otherwise
- $\text{cost} = \sum(\text{Buy-in cost})_i \cdot y_i + \sum(\text{cost/unit})_i \cdot x_i$
- $\text{payoff} = \sum(\text{payoff/unit})_i \cdot x_i$
- $\text{profit} = \text{payoff} - \text{cost}$

In addition, there are several constraints on the investments:

- You may not invest in both **2** and **5**.
- You may invest in **1** only if you invest in at least one of **2** and **3**.
- You must invest in at least two of **3**, **4**, and **5**.
- You may not invest more than the listed maximum number of units in each investment.

The constraints are represented in symbols as follows:

- $y_2 + y_5 \leq 1$
- $y_1 \leq y_2 + y_3$
- $y_3 + y_4 + y_5 \geq 2$
- $x_1 \leq 5$; $x_2 \leq 4$; $x_3 \leq 5$; $x_4 \leq 7$; $x_5 \leq 3$
- $\text{cost} \leq 125$

bintprog Formulation

To frame this problem as a binary integer programming problem, perform the following steps:

- 1 Represent each integer variable x_i by three binary integer variables $z_{i,j}$, $j = 1, \dots, 3$, as follows:

$$x_i = z_{i,1} + 2z_{i,2} + 4z_{i,3}, \quad i = 1, \dots, 5.$$

Three $z_{i,j}$ suffice to represent x_i , since each x_i is assumed to be 7 or less. And, since $x_5 \leq 3$, $z_{5,3} = 0$.

- 2 Combine the variables y and z into a single vector t as follows:

t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	t_{11}	t_{12}	t_{13}	t_{14}	t_{15}	t_{16}	t_{17}	t_{18}	t_{19}
y_1	y_2	y_3	y_4	y_5	$z_{1,1}$	$z_{1,2}$	$z_{1,3}$	$z_{2,1}$	$z_{2,2}$	$z_{2,3}$	$z_{3,1}$	$z_{3,2}$	$z_{3,3}$	$z_{4,1}$	$z_{4,2}$	$z_{4,3}$	$z_{5,1}$	$z_{5,2}$

- 3 Include the constraints $y_i = 0$ if and only if all the corresponding $z_{i,j} = 0$ as follows:

- $y_i \leq z_{i,1} + z_{i,2} + z_{i,3}$
- $z_{i,1} + 2z_{i,2} + 4z_{i,3} \leq y_i^* (\text{Max \# units}_i)$

These two inequalities enforce $y_i = 0$ if and only if all the $z_{i,j} = 0$, and they also enforce the maximum # units constraints.

- 4 As described in “Maximizing an Objective” on page 2-31, you find a maximum of an objective function by minimizing the negative of the objective function. So, to find the maximum profit, minimize the negative of the profit. The vector f that gives the negative of the profit in the form $f' * t$ is

$$f = [25, 35, 28, 20, 40, -10, -20, -40, -18, -36, -72, -11, -22, -44, \dots, -9, -18, -36, -10, -20]'$$

- The first five entries in f represent the buy-in costs; these are incurred if the corresponding $y_i = 1$.

- The next three entries, $f(6)$, $f(7)$, and $f(8)$, represent the negative of payoff minus cost per unit for investment 1, $-(\$15 - \$5)$, multiplied by 1, 2, and 4 respectively.
- The entries $f(9)$, $f(10)$, and $f(11)$ represent the corresponding quantities for investment 2: $-(\$25 - \$7)$, multiplied by 1, 2, and 4.
- $f(12)$, $f(13)$, and $f(14)$ correspond to investment 3
- $f(15)$, $f(16)$, and $f(17)$ correspond to investment 4
- $f(18)$ and $f(19)$ correspond to investment 5

5 Formulate all the constraints as inequalities of the form $A \cdot t \leq b$, as required in the `binprog` formulation “Binary Integer Programming Definition” on page 6-148.

The following matrix A represents the constraints, along with the vector b :

```
A = zeros(14,19);
A(1,1:19) = [25 35 28 20 40 5 10 20 7 14 28 ...
             6 12 24 4 8 16 8 16];
A(2,1) = 1; A(2,6) = -1; A(2,7) = -1; A(2,8) = -1;
A(3,2) = 1; A(3,9) = -1; A(3,10) = -1; A(3,11) = -1;
A(4,3) = 1; A(4,12) = -1; A(4,13) = -1; A(4,14) = -1;
A(5,4) = 1; A(5,15) = -1; A(5,16) = -1; A(5,17) = -1;
A(6,5) = 1; A(6,18) = -1; A(6,19) = -1;
A(7,1) = -5; A(7,6) = 1; A(7,7) = 2; A(7,8) = 4;
A(8,2) = -4; A(8,9) = 1; A(8,10) = 2; A(8,11) = 4;
A(9,3) = -5; A(9,12) = 1; A(9,13) = 2; A(9,14) = 4;
A(10,4) = -7; A(10,15) = 1; A(10,16) = 2; A(10,17) = 4;
A(11,5) = -3; A(11,18) = 1; A(11,19) = 2;
A(12,2) = 1; A(12,5) = 1;
A(13,1) = 1; A(13,2) = -1; A(13,3) = -1;
A(14,3) = -1; A(14,4) = -1; A(14,5) = -1;
b = [125 0 0 0 0 0 0 0 0 0 0 1 0 -2]';
```

- The first row of A represents the cost structure; no more than \$125 is available.
- Rows 2 through 6 represent $y_i \leq \sum_j z_{i,j}$, $i = 1, \dots, 5$.

- Rows 7 through 11 represent the maximum # units constraints. They also enforce $y_i = 1$ when $\sum_j z_{ij} > 0$.
- Rows 12, 13, and 14 represent the other constraints on investments.

bintprog Solution

bintprog solves the optimization problem as follows:

```
[t fval exitflag output] = bintprog(f,A,b);
Optimization terminated.
```

To examine the result, enter

```
t',fval
```

```
ans =
  Columns 1 through 10
    0    0    1    1    0    0    0    0    0    0
  Columns 11 through 19
    0    1    0    1    1    1    1    0    0

fval =
   -70
```

You can easily see that the only positive values of y are y_3 and y_4 . The values of x that correspond to these, x_3 and x_4 , are

```
t(12) + 2*t(13) + 4*t(14)
ans =
    5
```

```
t(15) + 2*t(16) + 4*t(17)
ans =
    7
```

In other words, to obtain the maximum profit, \$70, invest in 5 units of **3** and 7 units of **4**. By the way, this uses only

$$28 + (5*6) + 20 + (7*4) = 106$$

of the \$125 you had to invest, so there is \$19 left uninvested. You can also see this by checking the first constraint, $[A*t](1)$:

```
A(1,:)*t  
ans =  
    106
```

Mixed-Integer Linear Programming Algorithms

In this section...

“Mixed-Integer Linear Programming Definition” on page 6-157

“intlinprog Algorithm” on page 6-157

Mixed-Integer Linear Programming Definition

A mixed-integer linear program is a problem with

- Linear objective function, $f^T x$, where f is a column vector of constants, and x is the column vector of unknowns
- Bounds and linear constraints, but no nonlinear constraints (for definitions, see “Writing Constraints” on page 2-33)
- Restrictions on some components of x to have integer values

In mathematical terms, given vectors f , lb , and ub , matrices A and Aeq , corresponding vectors b and beq , and a set of indices `intcon`, find a vector x to solve

$$\min_x f^T x \text{ subject to } \begin{cases} x(\text{intcon}) \text{ are integers} \\ A \cdot x \leq b \\ Aeq \cdot x = beq \\ lb \leq x \leq ub. \end{cases}$$

intlinprog Algorithm

- “Algorithm Overview” on page 6-158
- “Linear Program Preprocessing” on page 6-158
- “Linear Programming” on page 6-159
- “Mixed-Integer Program Preprocessing” on page 6-159
- “Cut Generation” on page 6-160
- “Heuristics for Finding Feasible Solutions” on page 6-161

- “Branch and Bound” on page 6-161

Algorithm Overview

`intlinprog` uses this basic strategy to solve mixed-integer linear programs. `intlinprog` can solve the problem in any of the stages. If it solves the problem in a stage, `intlinprog` does not execute the later stages.

- 1 Reduce the problem size using “Linear Program Preprocessing” on page 6-158.
- 2 Solve an initial relaxed (noninteger) problem using “Linear Programming” on page 6-159.
- 3 Perform “Mixed-Integer Program Preprocessing” on page 6-159 to tighten the LP relaxation of the mixed-integer problem.
- 4 Try “Cut Generation” on page 6-160 to further tighten the LP relaxation of the mixed-integer problem.
- 5 Try to find integer-feasible solutions using heuristics.
- 6 Use a “Branch and Bound” on page 6-161 algorithm to search systematically for the optimal solution. This algorithm solves LP relaxations with restricted ranges of possible values of the integer variables. It attempts to generate a sequence of updated bounds on the optimal objective function value.

Linear Program Preprocessing

According to the “Mixed-Integer Linear Programming Definition” on page 6-157, there are matrices A and A_{eq} and corresponding vectors b and b_{eq} that encode a set of linear inequalities and linear equalities

$$A \cdot x \leq b$$
$$A_{eq} \cdot x = b_{eq}.$$

These linear constraints restrict the solution x .

Usually, it is possible to reduce the number of variables in the problem (the number of components of x), and reduce the number of linear constraints.

While performing these reductions can take time for the solver, they usually lower the overall time to solution, and can make larger problems solvable. The algorithms can make solution more numerically stable. Furthermore, these algorithms can sometimes detect an infeasible problem.

Preprocessing steps aim to eliminate redundant variables and constraints, improve the scaling of the model and sparsity of the constraint matrix, strengthen the bounds on variables, and detect the primal and dual infeasibility of the model.

For details, see Andersen and Andersen [1] and Mészáros and Suhl [4].

Linear Programming

The initial *relaxed* problem is the linear programming problem with the same objective and constraints as “Mixed-Integer Linear Programming Definition” on page 6-157, but no integer constraints. Call x_{LP} the solution to the relaxed problem, and x the solution to the original problem with integer constraints. Clearly,

$$f^T x_{LP} \leq f^T x,$$

because x_{LP} minimizes the same function but with fewer restrictions.

This initial relaxed LP (root node LP) and all generated LP relaxations during the branch-and-bound algorithm are solved using linear programming solution techniques.

Mixed-Integer Program Preprocessing

During mixed-integer program preprocessing, `intlinprog` analyzes the linear inequalities $A*x \leq b$ along with integrality restrictions to determine whether:

- The problem is infeasible.
- Some bounds can be tightened.
- Some inequalities are redundant, so can be ignored or removed.
- Some inequalities can be strengthened.
- Some integer variables can be fixed.

The `IPPreprocess` option lets you choose whether `intlinprog` takes several steps, takes all of them, or takes almost none of them.

The main goal of mixed-integer program preprocessing is to simplify ensuing branch-and-bound calculations. Preprocessing involves quickly preexamining and eliminating some of the futile subproblem candidates that branch-and-bound would otherwise analyze.

For details about integer preprocessing, see Savelsbergh [6].

Cut Generation

Cuts are additional linear inequality constraints that `intlinprog` adds to the problem. These inequalities attempt to restrict the feasible region of the LP relaxations so that their solution are closer to integers. You control the type of cuts that `intlinprog` uses with the `CutGeneration` option.

'basic' cuts include:

- Mixed-integer rounding cuts
- Gomory cuts
- Cliques cuts
- Cover cuts
- Flow cover cuts

'intermediate' cuts include all 'basic' cuts, plus:

- Simple lift-and-project cuts
- Simple pivot-and-reduce cuts
- Reduce-and-split cuts

'advanced' cuts include all 'intermediate' cuts except reduce-and-split cuts, plus:

- Strong Chvatal-Gomory cuts
- Zero-half cuts

Another option, `CutGenMaxIter`, specifies an upper bound on the number of times `intlinprog` iterates to generate cuts.

For details about cut generation algorithms (also called cutting plane methods), see Cornuéjols [2].

Heuristics for Finding Feasible Solutions

To get an upper bound on the objective function, the branch-and-bound procedure must find feasible points. A solution to an LP relaxation during branch-and-bound can be integer feasible, which can provide an improved upper bound to the original MILP. There are techniques for finding feasible points faster before and/or during branch-and-bound. These techniques are heuristic, meaning they are algorithms that can succeed, but can also fail. You set the `intlinprog` heuristics in the `Heuristics` option. The options are:

- `'rins'` — `intlinprog` searches the neighborhood of the current best integer feasible solution point (if available) to find a new and better solution. See Danna, Rothberg, and Le Pape [3].
- `'rss'` — `intlinprog` applies a hybrid procedure combining ideas from `'rins'` and local branching to search for integer feasible solutions.
- `'round'` — `intlinprog` takes the LP solution to the relaxed problem at a node. It rounds the integer components in a way that attempts to maintain feasibility.
- `'none'` — `intlinprog` does not search for a feasible point. It simply takes any feasible point it encounters in its branch-and-bound search.

Branch and Bound

The branch-and-bound method constructs a sequence of subproblems that attempt to converge to a solution of the MILP. The subproblems give a sequence of upper and lower bounds on the solution $f^T x$. The first upper bound is any feasible solution, and the first lower bound is the solution to the relaxed problem. For a discussion of the upper bound, see “Heuristics for Finding Feasible Solutions” on page 6-161.

As explained in “Linear Programming” on page 6-159, any solution to the linear programming relaxed problem has a lower objective function value than the solution to the MILP. Also, any feasible point x_{feas} satisfies

$$f^T x_{\text{feas}} \geq f^T x,$$

because $f^T x$ is the minimum among all feasible points.

In this context, a *node* is an LP with the same objective function, bounds, and linear constraints as the original problem, but without integer constraints, and with particular changes to the linear constraints or bounds. The *root node* is the original problem with no integer constraints and no changes to the linear constraints or bounds, meaning the root node is the initial relaxed LP.

From the starting bounds, the branch-and-bound method constructs new subproblems by branching from the root node. The branching step is taken heuristically, according to one of several rules. Each rule is based on the idea of splitting a problem by restricting one variable to be less than or equal to an integer J , or greater than or equal to $J+1$. These two subproblems arise when an entry in x_{LP} , corresponding to an integer specified in `intcon`, is not an integer. Here, x_{LP} is the solution to a relaxed problem. Take J as the floor of the variable (rounded down), and $J+1$ as the ceiling (rounded up). The resulting two problems have solutions that are larger than or equal to $f^T x_{\text{LP}}$, because they have more restrictions. Therefore, this procedure potentially raises the lower bound.

The performance of the branch-and-bound method depends on the rule for choosing which variable to split (the branching rule). The algorithm uses these rules, which you can set in the `BranchingRule` option:

- 'maxpscost' — Choose the fractional variable with maximal *pseudocost*.

Pseudocost

The pseudocost of a variable i is based on empirical estimates of the change in the lower bound when i has been chosen as the branching variable, combined with the fractional part of the i component of the current point x . The fractional part p is in two pieces, the lower part and the upper part:

$$\begin{aligned} p_i^- &= x(i) - x(i) \\ p_i^+ &= 1 - p_i^- \end{aligned}$$

Let x_i^- be the solution of the linear program restricted to have $x(i) \leq x(i)$, and let the change in objective function be denoted

$$\Delta_i^- = f^T x_i^- - f^T x.$$

Similarly, Δ_i^+ is the change in objective function when the problem is restricted to have $x(i) \geq x(i)$.

The objective gain per unit change in variable x_i is

$$d_i^- = \frac{\Delta_i^-}{p_i^-} \text{ or } d_i^+ = \frac{\Delta_i^+}{p_i^+}.$$

Let s_i^- and s_i^+ be the empirical averages of d_i^- and d_i^+ during the branch-and-bound algorithm up to this point. The empirical values are initialized to the absolute value of the objective coefficient $f(i)$ for the terms before there are any observations. Then the 'maxpscost' rule is to branch on a node i that maximizes, for some positive weights w^+ and w^- , the quantity

$$w^- * p_i^- * s_i^- + w^+ * p_i^+ * s_i^+.$$

Roughly speaking, this rule chooses a coefficient that is likely to increase the lower bound maximally.

- 'mostfractional' — Choose the variable with fractional part closest to 1/2.
- 'maxfun' — Choose the variable with maximal corresponding absolute value in the objective vector f .

After the algorithm branches, there are two new nodes to explore. The algorithm chooses which node to explore among all that are available using one of these rules:

- 'minobj' — Choose the node that has the lowest objective function value.
- 'mininfeas' — Choose the node with the minimal sum of integer infeasibilities. This means for every integer-infeasible component $x(i)$ in the node, add up the smaller of p_i^- and p_i^+ , where

$$\begin{aligned} p_i^- &= x(i) - x(i) \\ p_i^+ &= 1 - p_i^- \end{aligned}$$

- 'simplebestproj' — Choose the node with the *best projection*.

Best Projection

Let x_B denote the best integer-feasible point found so far, x_R denote the LP relaxed solution at the root node, and x denote the node we examine. Let $in(x)$ denote the sum of integer infeasibilities at the node x (see 'mininfeas'). The best projection rule is to minimize

$$f^T x + \frac{f^T x_B - f^T x_R}{in(x_R)} in(x).$$

If there is no integer-feasible point found so far, set $f^T x_B = 0$.

The branch-and-bound procedure continues, systematically generating subproblems to analyze and discarding the ones that won't improve an upper or lower bound on the objective, until one of these stopping criteria is met:

- The algorithm exceeds the MaxTime option.
- The difference between the lower and upper bounds on the objective function is less than the TolGapAbs or TolGapRel tolerances.
- The number of explored nodes exceeds the MaxNodes option.
- The number of integer feasible points exceeds the MaxNumFeasPoints option.

For details about the branch-and-bound procedure, see Nemhauser and Wolsey [5] and Wolsey [7].

References

- [1] Andersen, E. D., and Andersen, K. D. *Presolving in linear programming*. Mathematical Programming 71, pp. 221–245, 1995.
- [2] Cornuéjols, G. *Valid inequalities for mixed integer linear programs*. Mathematical Programming B, Vol. 112, pp. 3–44, 2008.

- [3] Danna, E., Rothberg, E., Le Pape, C. *Exploring relaxation induced neighborhoods to improve MIP solutions*. Mathematical Programming, Vol. 102, issue 1, pp. 71–90, 2005.
- [4] Mészáros C., and Suhl, U. H. *Advanced preprocessing techniques for linear and quadratic programming*. OR Spectrum, 25(4), pp. 575–595, 2003.
- [5] Nemhauser, G. L. and Wolsey, L. A. *Integer and Combinatorial Optimization*. Wiley-Interscience, New York, 1999.
- [6] Savelsbergh, M. W. P. *Preprocessing and Probing Techniques for Mixed Integer Programming Problems*. ORSA J. Computing, Vol. 6, No. 4, pp. 445–454, 1994.
- [7] Wolsey, L. A. *Integer Programming*. Wiley-Interscience, New York, 1998.

Tuning Integer Linear Programming

In this section...
“Change Options to Improve the Solution Process” on page 6-166
“Some “Integer” Solutions Are Not Integers” on page 6-167
“Large Components Not Integer Valued” on page 6-168
“Large Coefficients Disallowed” on page 6-168

Change Options to Improve the Solution Process

Note Often, you can change the formulation of a MILP to make it more easily solvable. For suggestions on how to change your formulation, see Williams [1].

After you run `intlinprog` once, you might want to change some options and rerun it. The changes you might want to see include:

- Lower run time
- Lower final objective function value (a better solution)
- Smaller final gap
- More or different feasible points

Here are general recommendations for option changes that are most likely to help the solution process. Try the suggestions in this order:

- 1** For a faster and more accurate solution, increase the `CutGenMaxIter` option from its default 10 to a higher number such as 25. This can speed up the solution, but can also slow it.
- 2** For a faster and more accurate solution, change the `CutGeneration` option to 'intermediate' or 'advanced'. This can speed up the solution, but can use much more memory, and can slow the solution.

- 3** For a faster and more accurate solution, change the `IPPreprocess` option to `'advanced'`. This can have a large effect on the solution process, either beneficial or not.
- 4** For a faster and more accurate solution, change the `RootLPAlgorithm` option to `'primal-simplex'`. Usually this change is not beneficial, but occasionally it can be.
- 5** To try to find more or better feasible points, increase the `HeuristicsMaxNodes` option from its default 50 to a higher number such as 100.
- 6** To try to find more or better feasible points, change the `Heuristics` option to either `'round'` or `'rins'`.
- 7** To attempt to stop the solver more quickly, change the `TolGapRel` option to a higher value than the default `1e-4`. Similarly, to attempt to obtain a more accurate answer, change the `TolGapRel` option to a lower value. These changes do not always improve results.
- 8** For a more accurate solution, decrease the `RelObjThreshold` option from its default `1e-4` to a smaller positive value such as `1e-6`. This change can cause `intlinprog` to take more time to solve the problem, and to find more integer feasible points during its solution process.

Some “Integer” Solutions Are Not Integers

Often, some supposedly integer-valued components of the solution `x(intcon)` are not precisely integers. `intlinprog` considers as integers all solution values within the `TolInteger` tolerance of an integer.

To round all supposed integers to be precisely integers, use the `round` function.

```
x(intcon) = round(x(intcon));
```

Caution Rounding can cause solutions to become infeasible. Check feasibility after rounding:

```
max(A*x - b) % see if entries are not too positive, so have small infeasibi
max(abs(Aeq*x - beq)) % see if entries are near enough to zero
max(x - ub) % positive entries are violated bounds
max(lb - x) % positive entries are violated bounds
```

Large Components Not Integer Valued

`intlinprog` does not enforce that solution components be integer valued when their absolute values exceed $2.1e9$. When your solution has such components, `intlinprog` warns you. If you receive this warning, check the solution to see whether supposedly integer-valued components of the solution are close to integers.

Large Coefficients Disallowed

`intlinprog` does not allow components of the problem, such as coefficients in f , A , or ub , to exceed $1e15$ in absolute value. If you try to run `intlinprog` with such a problem, `intlinprog` issues an error.

If you get this error, sometimes you can scale the problem to have smaller coefficients:

- For coefficients in f that are too large, try multiplying f by a small positive scaling factor.
- For constraint coefficients that are too large, try multiplying all bounds and constraint matrices by the same small positive scaling factor.

References

[1] Williams, H. Paul. *Model Building in Mathematical Programming*. Wiley, 2013.

Mixed-Integer Linear Programming Basics

This example shows how to solve a mixed-integer linear program. The example is not complex, but it shows typical steps in formulating a problem in the syntax for `intlinprog`.

Problem description

You want to blend a variety of steels with various chemical compositions to obtain 25 tons of steel with a specific chemical composition. The result should have 5% carbon and 5% molybdenum by weight, meaning $25 \text{ tons} \times 5\% = 1.25$ tons of carbon and 1.25 tons of molybdenum. The objective is to minimize the cost for blending the steel.

This problem is taken from Carl-Henrik Westerberg, Bengt Bjorklund and Eskil Hultman, “*An Application of Mixed Integer Programming in a Swedish Steel Mill.*” *Interfaces* February 1977 Vol. 7, No. 2 pp. 39–43, whose abstract is at <http://interfaces.journal.informs.org/content/7/2/39.abstract>.

Four ingots of steel are available for purchase. Only one of each ingot is available.

Ingot	Weight (tons)	%Carbon	%Molybdenum	Cost/ton
1	5	5	3	\$350
2	3	4	3	\$330
3	4	5	4	\$310
4	6	3	4	\$280

Three grades of alloy steel are available for purchase, and one grade of scrap steel. Alloy and scrap steels can be purchased in fractional amounts.

Alloy	%Carbon	%Molybdenum	Cost/ton
1	8	6	\$500
2	7	7	\$450

Alloy	%Carbon	%Molybdenum	Cost/ton
3	6	8	\$400
Scrap	3	9	\$100

To formulate the problem, first decide on the control variables. Take variable $x(1) = 1$ to mean you purchase ingot 1, and $x(1) = 0$ to mean you do not purchase the ingot. Similarly, variables $x(2)$ through $x(4)$ are binary variables indicating that you purchase ingots 2 through 4.

Variables $x(5)$ through $x(7)$ are the quantities of alloys 1, 2, and 3 you purchase, and $x(8)$ is the quantity of scrap steel you purchase.

MATLAB formulation

Formulate the problem by specifying the inputs for `intlinprog`. The relevant `intlinprog` syntax is as follows.

```
[x,fval] = intlinprog(f,intcon,A,b,Aeq,beq,lb,ub)
```

Create the inputs for `intlinprog` from first (f) through last (ub).

f is the vector of cost coefficients. The coefficients representing the costs of ingots are the ingot weights times their cost per ton.

```
f = [350*5,330*3,310*4,280*6,500,450,400,100];
```

The integer variables are the first four.

```
intcon = 1:4;
```

Tip To specify binary variables, set the variables to be integers in `intcon`, and give them a lower bound of 0 and an upper bound of 1.

There are no linear inequality constraints, so A and b are empty [].

There are three equality constraints. The first is that the total weight is 25 tons.

$$5*x(1) + 3*x(2) + 4*x(3) + 6*x(4) + x(5) + x(6) + x(7) + x(8) = 25.$$

The second constraint is that the weight of carbon is 5% of 25 tons, or 1.25 tons.

$$5*0.05*x(1) + 3*0.04*x(2) + 4*0.05*x(3) + 6*0.03*x(4) + 0.08*x(5) + 0.07*x(6) + 0.06*x(7) + 0.03*x(8) = 1.25.$$

The third constraint is that the weight of molybdenum is 1.25 tons.

$$5*0.03*x(1) + 3*0.03*x(2) + 4*0.04*x(3) + 6*0.04*x(4) + 0.06*x(5) + 0.07*x(6) + 0.08*x(7) + 0.09*x(8) = 1.25.$$

In matrix form, $Aeq*x = beq$, where

```
Aeq = [5,3,4,6,1,1,1,1;
       5*0.05,3*0.04,4*0.05,6*0.03,0.08,0.07,0.06,0.03;
       5*0.03,3*0.03,4*0.04,6*0.04,0.06,0.07,0.08,0.09];
beq = [25;1.25;1.25];
```

Each variable is bounded below by zero. The integer variables are bounded above by one.

```
lb = zeros(8,1);
ub = ones(8,1);
ub(5:end) = Inf; % No upper bound on noninteger variables
```

Solve the problem

Now that you have all the inputs, call the solver.

```
[x,fval] = intlinprog(f,intcon,[],[],Aeq,beq,lb,ub);
```

View the solution.

```
x,fval
```

```
x =
```

```
1.0000
1.0000
```

```
0
1.0000
7.2500
0
0.2500
3.5000
```

```
fval =
```

```
8.4950e+03
```

The optimal purchase costs \$8,495. Buy ingots **1**, **2**, and **4**, but not **3**, and buy 7.25 tons of alloy **1**, 0.25 ton of alloy **3**, and 3.5 tons of scrap steel.

Set `intcon = []` to see the effect of solving the problem without integer constraints. The solution is different, and is not sensible, because you cannot purchase a fraction of an ingot.

Factory, Warehouse, Sales Allocation Model

This example shows how to set up and solve a mixed-integer linear programming problem. The problem is to find the optimal production and distribution levels among a set of factories, warehouses, and sales outlets.

The example first generates random locations for factories, warehouses, and sales outlets. Feel free to modify the scaling parameter α , which scales both the size of the grid in which the production and distribution facilities reside, but also scales the number of these facilities so that the density of facilities of each type per grid area is independent of α .

Facility Locations

For a given value of the scaling parameter α , suppose that there are the following:

- factories
- warehouses
- sales outlets

These facilities are on separate integer grid points between 1 and α in the x and y directions. In order that the facilities have separate locations, you require that $\alpha > 1$. In this example, take $\alpha = 100$, $\alpha = 100$, and $\alpha = 100$.

Production and Distribution

There are P products made by the factories. Take $P = 10$.

The demand for each product p in a sales outlet s is d_{ps} . The demand is the quantity that can be sold in a time interval. One constraint on the model is that the demand is met, meaning the system produces and distributes exactly the quantities in the demand.

There are capacity constraints on each factory and each warehouse.

- The production of product p at factory f is less than C_f .
- The capacity of warehouse w is C_w .

- The amount of product that can be transported from warehouse to a sales outlet in the time interval is less than , where is the turnover rate of product .

Suppose that each sales outlet receives its supplies from just one warehouse. Part of the problem is to determine the cheapest mapping of sales outlets to warehouses.

Costs

The cost of transporting products from factory to warehouse, and from warehouse to sales outlet, depends on the distance between the facilities, and on the particular product. If is the distance between facilities and , then the cost of shipping a product between these facilities is the distance times the transportation cost :

The distance in this example is the grid distance, also known as the distance. It is the sum of the absolute difference in coordinates and coordinates.

The cost of making a unit of product in factory is .

Optimization Problem

Given a set of facility locations, and the demands and capacity constraints, find:

- A production level of each product at each factory
- A distribution schedule for products from factories to warehouses
- A distribution schedule for products from warehouses to sales outlets

These quantities must ensure that demand is satisfied and total cost is minimized. Also, each sales outlet is required to receive all its products from exactly one warehouse.

Variables and Equations for the Optimization Problem

The control variables, meaning the ones you can change in the optimization, are

- = the amount of product that is transported from factory to warehouse
- = a binary variable taking value 1 when sales outlet is associated with warehouse

The objective function to minimize is

The constraints are

(capacity of factory).

(demand is met).

(capacity of warehouse).

(each sales outlet associates to one warehouse).

(nonnegative production).

(binary).

The variables and appear in the objective and constraint functions linearly. Because is restricted to integer values, the problem is a mixed-integer linear program (MILP).

Generate a Random Problem: Facility Locations

Set the values of the , , , and parameters, and generate the facility locations.

```
rng default % for reproducibility
N = 20; % N from 10 to 30 seems to work. Choose large values with caution.
N2 = N*N;
f = 0.05; % density of factories
w = 0.05; % density of warehouses
```

```
s = 0.1; % density of sales outlets

F = floor(f*N2); % number of factories
W = floor(w*N2); % number of warehouses
S = floor(s*N2); % number of sales outlets

xyloc = randperm(N2,F+W+S); % unique locations of facilities
[xloc,yloc] = ind2sub([N N],xyloc);
```

Of course, it is not realistic to take random locations for facilities. This example is intended to show solution techniques, not how to generate good facility locations.

Plot the facilities. Facilities 1 through F are factories, F+1 through F+W are warehouses, and F+W+1 through F+W+S are sales outlets.

```
plot(xloc(1:F),yloc(1:F),'rs',xloc(F+1:F+W),yloc(F+1:F+W),'k*','...
      xloc(F+W+1:F+W+S),yloc(F+W+1:F+W+S),'bo')
legend('Factory','Warehouse','Sales outlet','Location','EastOutside')
xlim([0 N+1]);ylim([0 N+1])
```

Generate Random Capacities, Costs, and Demands

Generate random production costs, capacities, turnover rates, and demands.

```
P = 20; % 20 products

% Production costs between 20 and 100
pcost = 80*rand(F,P) + 20;

% Production capacity between 500 and 1500 for each product/factory
pcap = 1000*rand(F,P) + 500;

% Warehouse capacity between P*400 and P*800 for each product/warehouse
wcap = P*400*rand(W,P) + P*400;

% Product turnover rate between 1 and 3 for each product
turn = 2*rand(1,P) + 1;
```

```

% Product transport cost per distance between 5 and 10 for each product
tcost = 5*rand(1,P) + 5;

% Product demand by sales outlet between 200 and 500 for each
% product/outlet
d = 300*rand(S,P) + 200;

```

These random demands and capacities can lead to infeasible problems. In other words, sometimes the demand exceeds the production and warehouse capacity constraints. If you alter some parameters and get an infeasible problem, during solution you will get an exitflag of -2.

Generate Objective and Constraint Matrices and Vectors

The objective function vector `obj` in `intlincon` consists of the coefficients of the variables `x` and `w`. So there are naturally $P \times F \times W + S \times W$ coefficients in `obj`.

One way to generate the coefficients is to begin with a P -by- F -by- W array `obj1` for the coefficients, and an S -by- W array `obj2` for the coefficients. Then convert these arrays to two vectors and combine them into `obj` by calling

```

obj = [obj1(:);obj2(:)];

obj1 = zeros(P,F,W); % Allocate arrays
obj2 = zeros(S,W);

```

Throughout the generation of objective and constraint vectors and matrices, we generate the `obj1` array or the `obj2` array, and then convert the result to a vector.

To begin generating the inputs, generate the distance arrays `distfw(i,j)` and `distsw(i,j)`.

```

distfw = zeros(F,W); % Allocate matrix for factory-warehouse distances
for ii = 1:F
    for jj = 1:W
        distfw(ii,jj) = abs(xloc(ii) - xloc(F + jj)) + abs(yloc(ii) ...
            - yloc(F + jj));
    end
end

```

```
distsw = zeros(S,W); % Allocate matrix for sales outlet-warehouse distances
for ii = 1:S
    for jj = 1:W
        distsw(ii,jj) = abs(xloc(F + W + ii) - xloc(F + jj)) ...
            + abs(yloc(F + W + ii) - yloc(F + jj));
    end
end
```

Generate the entries of obj1 and obj2.

```
for ii = 1:P
    for jj = 1:F
        for kk = 1:W
            obj1(ii,jj,kk) = pcost(jj,ii) + tcost(ii)*distfw(jj,kk);
        end
    end
end
```

```
for ii = 1:S
    for jj = 1:W
        obj2(ii,jj) = distsw(ii,jj)*sum(d(ii,:).*tcost);
    end
end
```

Combine the entries into one vector.

```
obj = [obj1(:);obj2(:)]; % obj is the objective function vector
```

Now create the constraint matrices.

The width of each linear constraint matrix is the length of the obj vector.

```
matwid = length(obj);
```

There are two types of linear inequalities: the production capacity constraints, and the warehouse capacity constraints.

There are $P \times F$ production capacity constraints, and W warehouse capacity constraints. The constraint matrices are quite sparse, on the order of 1% nonzero, so save memory by using sparse matrices.

```
Aineq = spalloc(P*F + W,matwid,P*F*W + S*W); % Allocate sparse Aeq
bineq = zeros(P*F + W,1); % Allocate bineq as full
```

```
% Zero matrices of convenient sizes:
```

```
clearer1 = zeros(size(obj1));
clearer12 = clearer1(:);
clearer2 = zeros(size(obj2));
clearer22 = clearer2(:);
```

```
% First the production capacity constraints
```

```
counter = 1;
```

```
for ii = 1:F
```

```
    for jj = 1:P
```

```
        xtemp = clearer1;
```

```
        xtemp(jj,ii,:) = 1; % Sum over warehouses for each product and fact
```

```
        xtemp = sparse([xtemp(:);clearer22]); % Convert to sparse
```

```
        Aineq(counter,:) = xtemp'; % Fill in the row
```

```
        bineq(counter) = pcap(ii,jj);
```

```
        counter = counter + 1;
```

```
    end
```

```
end
```

```
% Now the warehouse capacity constraints
```

```
vj = zeros(S,1); % The multipliers
```

```
for jj = 1:S
```

```
    vj(jj) = sum(d(jj,)./turn); % A sum of P elements
```

```
end
```

```
for ii = 1:W
```

```
    xtemp = clearer2;
```

```
    xtemp(:,ii) = vj;
```

```
    xtemp = sparse([clearer12;xtemp(:)]); % Convert to sparse
```

```
    Aineq(counter,:) = xtemp'; % Fill in the row
```

```
    bineq(counter) = wcap(ii);
```

```
    counter = counter + 1;
```

```
end
```

There are two types of linear equality constraints: the constraint that demand is met, and the constraint that each sales outlet corresponds to one warehouse.

```
Aeq = spalloc(P*W + S,matwid,P*W*(F+S) + S*W); % Allocate as sparse
beq = zeros(P*W + S,1); % Allocate vectors as full

counter = 1;
% Demand is satisfied:
for ii = 1:P
    for jj = 1:W
        xtemp = clearer1;
        xtemp(ii,:,jj) = 1;
        xtemp2 = clearer2;
        xtemp2(:,jj) = -d(:,ii);
        xtemp = sparse([xtemp(:);xtemp2(:)]'); % Change to sparse row
        Aeq(counter,:) = xtemp; % Fill in row
        counter = counter + 1;
    end
end

% Only one warehouse for each sales outlet:
for ii = 1:S
    xtemp = clearer2;
    xtemp(ii,:) = 1;
    xtemp = sparse([clearer12;xtemp(:)]'); % Change to sparse row
    Aeq(counter,:) = xtemp; % Fill in row
    beq(counter) = 1;
    counter = counter + 1;
end
```

Bound Constraints and Integer Variables

The integer variables are those from `length(obj1) + 1` to the end.

```
intcon = P*F*W+1:length(obj);
```

The upper bounds are from `length(obj1) + 1` to the end also.

```
lb = zeros(length(obj),1);
```

```
ub = Inf(length(obj),1);
ub(P*F*W+1:end) = 1;
```

Turn off iterative display so that you don't get hundreds of lines of output.

```
opts = optimoptions('intlinprog','Display','off');
```

Solve the Problem

You generated all the solver inputs. Call the solver to find the solution.

```
[solution,fval,exitflag,output] = intlinprog(obj,intcon,...
                                             Aineq,bineq,Aeq,beq,lb,ub,opts);
if exitflag < 0 % If you changed some parameters and got an infeasible prob
    disp('The problem with these parameters is infeasible. No solution.')
    break % Stop the script because there is nothing to examine
end
```

Examine the Solution

The solution is feasible, to within the given tolerances.

```
exitflag
infeas1 = max(Aineq*solution - bineq)
infeas2 = norm(Aeq*solution - beq,Inf)
```

```
exitflag =
```

```
1
```

```
infeas1 =
```

```
1.1369e-12
```

```
infeas2 =
```

```
7.9581e-13
```

Check that the integer components are really integers, or are close enough that it is reasonable to round them. To understand why these variables might not be exactly integers, see the documentation.

```
diffint = norm(solution(intcon) - round(solution(intcon)),Inf)
```

```
diffint =
```

```
4.4409e-16
```

Some integer variables are not exactly integers, but all are very close. So round the integer variables.

```
solution(intcon) = round(solution(intcon));
```

Check the feasibility of the rounded solution, and the change in objective function value.

```
infeas1 = max(Aineq*solution - bineq)
infeas2 = norm(Aeq*solution - beq,Inf)
diffrounding = norm(fval - obj(:)'*solution,Inf)
```

```
infeas1 =
```

```
1.1369e-12
```

```
infeas2 =
```

```
1.0232e-12
```

```
diffrounding =
```

```
4.8429e-08
```


Rounding the solution did not appreciably change its feasibility.

You can examine the solution most easily by reshaping it back to its original dimensions.

```
solution1 = solution(1:P*F*W); % The continuous variables
solution2 = solution(intcon); % The integer variables
solution1 = reshape(solution1,P,F,W);
solution2 = reshape(solution2,S,W);
```

For example, how many sales outlets are associated with each warehouse? Notice that, in this case, some warehouses have 0 associated outlets, meaning the warehouses are not in use in the optimal solution.

```
outlets = sum(solution2,1) % Sum over the sales outlets
```

```
outlets =
```

```
Columns 1 through 13
```

```
    2    0    2    3    3    2    2    3    2    3    2    1
```

```
Columns 14 through 20
```

```
    1    4    0    3    3    2    0
```

Plot the connection between each sales outlet and its warehouse.

```
hold on
```

```
for ii = 1:S
```

```
    jj = find(solution2(ii,:)); % Index of warehouse associated with ii
```

```
    xsales = xloc(F+W+ii); ysales = yloc(F+W+ii);
```

```
    xwarehouse = xloc(F+jj); ywarehouse = yloc(F+jj);
```

```
    if rand(1) < .5 % Draw y direction first half the time
```

```
        plot([xsales,xsales,xwarehouse],[ysales,ywarehouse,ywarehouse], 'g-'
```

```
    else % Draw x direction first the rest of the time
```

```
        plot([xsales,xwarehouse,xwarehouse],[ysales,ysales,ywarehouse], 'g- -')
    end
end
hold off

title('Mapping of sales outlets to warehouses')
```

The two black * with no green lines represent the unused warehouses.

Travelling Salesman Problem

This example shows how to use binary integer programming to solve the classic travelling salesman problem. This problem involves finding the shortest closed tour (path) through a set of stops (cities). In this case there are 200 stops, but you can easily change the `nStops` variable to get a different problem size. You'll solve the initial problem and see that the solution has subtours. This means the optimal solution found doesn't give one continuous path through all the points, but instead has several disconnected loops. You'll then use an iterative process of determining the subtours, adding constraints, and rerunning the optimization until the subtours are eliminated.

Draw the Map and Stops

Generate random stops inside a crude polygonal representation of the continental U.S.

```
figure;
```

```
load('usborder.mat','x','y','xx','yy');
rng(3,'twister') % makes a plot with stops in Maine & Florida, and is repro
nStops = 200; % you can use any number, but the problem size scales as N^2
stopsLon = zeros(nStops,1); % allocate x-coordinates of nStops
stopsLat = stopsLon; % allocate y-coordinates
n = 1;
while (n <= nStops)
    xp = rand*1.5;
    yp = rand;
    if inpolygon(xp,yp,xx,yy) % test if inside the border
        stopsLon(n) = xp;
        stopsLat(n) = yp;
        n = n+1;
    end
end
plot(x,y,'Color','red'); % draw the outside border
hold on
% Add the stops to the map
plot(stopsLon,stopsLat,'*b')
hold off
```

Problem Formulation

Formulate the travelling salesman problem for integer linear programming as follows:

- Generate all possible trips, meaning all distinct pairs of stops.
- Calculate the distance for each trip.
- The cost function to minimize is the sum of the trip distances for each trip in the tour.
- The decision variables are binary, and associated with each trip, where each 1 represents a trip that exists on the tour, and each 0 represents a trip that is not on the tour.
- To ensure that the tour includes every stop, include the linear constraint that each stop is on exactly two trips. This means one arrival and one departure from the stop.

Calculate Distances Between Points

Because there are 200 stops, there are 19,900 trips, meaning 19,900 binary variables (# variables = 200 choose 2).

Generate all the trips, meaning all pairs of stops.

```
idxs = nchoosek(1:nStops,2);
```

Calculate all the trip distances, assuming that the earth is flat in order to use the Pythagorean rule.

```
dist = hypot(stopsLat(idxs(:,1)) - stopsLat(idxs(:,2)), ...  
            stopsLon(idxs(:,1)) - stopsLon(idxs(:,2)));  
lendist = length(dist);
```

With this definition of the `dist` vector, the length of a tour is

```
dist'*x
```

where x is the binary solution vector. This is the distance of a tour that you try to minimize.

Equality Constraints

The problem has two types of equality constraints. The first enforces that there must be 200 trips total. The second enforces that each stop must have two trips attached to it (there must be a trip to each stop and a trip departing each stop).

Specify the first type of equality constraint, that you must have $nStops$ trips, in the form $Aeq*x = beq$.

```
Aeq = spones(1:length(idxs)); % Adds up the number of trips
beq = nStops;
```

To specify the second type of equality constraint, that there needs to be two trips attached to each stop, extend the Aeq matrix as sparse.

```
Aeq = [Aeq; spalloc(nStops,length(idxs),nStops*(nStops-1))]; % allocate a sp
for ii = 1:nStops
    whichIdxs = (idxs == ii); % find the trips that include stop ii
    whichIdxs = sparse(sum(whichIdxs,2)); % include trips where ii is at ei
    Aeq(ii+1,:) = whichIdxs'; % include in the constraint matrix
end
beq = [beq; 2*ones(nStops,1)];
```

Binary Bounds

All decision variables are binary. Now, set the `intcon` argument to the number of decision variables, put a lower bound of 0 on each, and an upper bound of 1.

```
intcon = 1:lendist;
lb = zeros(lendist,1);
ub = ones(lendist,1);
```

Optimize Using `intlinprog`

The problem is ready to be solved. Call the solver.

```
opts = optimoptions('intlinprog','Display','off');  
[xopt,costopt,exitflag,output] = intlinprog(dist,intcon,[],[],Aeq,beq,lb,ub,opts);
```

Visualize the Solution

```
hold on  
segments = find(xopt); % Get indices of lines on optimal path  
lh = zeros(nStops,1); % Use to store handles to lines on plot  
lh = updateSalesmanPlot(lh,xopt,idxs,stopsLon,stopsLat);  
title('Solution with Subtours');
```

As can be seen on the map, the solution has several subtours. The constraints specified so far do not prevent these subtours from happening. In order to prevent any possible subtour from happening, you would need an incredibly large number of inequality constraints.

Subtour Constraints

Because you can't add all of the subtour constraints, take an iterative approach. Detect the subtours in the current solution, then add inequality constraints to prevent those particular subtours from happening. By doing this, you find a suitable tour in a few iterations.

Eliminate subtours with inequality constraints. An example of how this works is if you have five points in a subtour, then you have five lines connecting those points to create the subtour. Eliminate this subtour by implementing an inequality constraint to say there must be less than or equal to four lines between these five points.

Even more, find all lines between these five points, and constrain the solution not to have more than four of these lines present. This is a correct constraint because if five or more of the lines existed in a solution, then the solution would have a subtour (a graph with nodes and edges always contains a cycle).

The `detectSubtours` function analyzes the solution and returns a cell array of vectors. Each vector in the cell array contains the stops involved in that particular subtour.

```
tours = detectSubtours(xopt,idxs);
numtours = length(tours); % number of subtours
fprintf('# of subtours: %d\n',numtours);
```

```
# of subtours: 28
```

Include the linear inequality constraints to eliminate subtours, and repeatedly call the solver, until just one subtour remains.

```
A = spalloc(0,lendist,0); % Allocate a sparse linear inequality constraint
b = [];
while numtours > 1 % repeat until there is just one subtour
    % Add the subtour constraints
    b = [b;zeros(numtours,1)]; % allocate b
    A = [A;spalloc(numtours,lendist,nStops)]; % a guess at how many nonzero
    for ii = 1:numtours
        rowIdx = size(A,1)+1; % Counter for indexing
        subTourIdx = tours{ii}; % Extract the current subtour
        % The next lines find all of the variables associated with the
        % particular subtour, then add an inequality constraint to prohibit
        % that subtour and all subtours that use those stops.
        variations = nchoosek(1:length(subTourIdx),2);
        for jj = 1:length(variations)
            whichVar = (sum(idxs==subTourIdx(variations(jj,1)),2)) & ...
                (sum(idxs==subTourIdx(variations(jj,2)),2));
            A(rowIdx,whichVar) = 1;
        end
        b(rowIdx) = length(subTourIdx)-1; % One less trip than subtour stops
    end

    % Try to optimize again
    [xopt,costopt,exitflag,output] = intlinprog(dist,intcon,A,b,Aeq,beq,lb,

    % Visualize result
    lh = updateSalesmanPlot(lh,xopt,idxs,stopsLon,stopsLat);
```

```
        % How many subtours this time?
        tours = detectSubtours(xopt,idxs);
        numtours = length(tours); % number of subtours
        fprintf('# of subtours: %d\n',numtours);
    end

    title('Solution with Subtours Eliminated');
    hold off

    # of subtours: 16
    # of subtours: 10
    # of subtours: 5
    # of subtours: 3
    # of subtours: 4
    # of subtours: 1
```

Solution Quality

The solution represents a feasible tour, because it is a single closed loop. But is it a minimal-cost tour? One way to find out is to examine the output structure.

output

```
output =

    relativegap: 0.0100
    absolutegap: 0.0021
    numfeaspoints: 5
    numnodes: 1226
    constrviolation: 2.0606e-13
    message: 'Optimal solution found.'

Intlinprog stopped because ...'
```


The gap between the internally calculated upper and lower bounds on the solution's objective function, reported in the output structure, is 0. So the solution is optimal.

Solve Sudoku Puzzles Via Integer Programming

This example shows how to solve a Sudoku puzzle using binary integer programming.

You probably have seen Sudoku puzzles. A puzzle is to fill a 9-by-9 grid with integers from 1 through 9 so that each integer appears only once in each row, column, and major 3-by-3 square. The grid is partially populated with clues, and your task is to fill in the rest of the grid.

Initial Puzzle

Here is a data matrix B of clues. The first row, $B(1,2,2)$, means row 1, column 2 has a clue 2. The second row, $B(1,5,3)$, means row 1, column 5 has a clue 3. Here is the entire matrix B .

```
B = [1,2,2;  
     1,5,3;  
     1,8,4;  
     2,1,6;  
     2,9,3;  
     3,3,4;  
     3,7,5;  
     4,4,8;  
     4,6,6;  
     5,1,8;  
     5,5,1;  
     5,9,6;  
     6,4,7;  
     6,6,5;  
     7,3,7;  
     7,7,6;  
     8,1,4;  
     8,9,8;  
     9,2,3;  
     9,5,4;  
     9,8,2];
```

```
drawSudoku(B) % For the listing of this program, see the end of this example
```

This puzzle, and an alternative MATLAB® solution technique, was featured in Cleve's Corner in 2009.

There are many approaches to solving Sudoku puzzles manually, as well as many programmatic approaches. This example shows a straightforward approach using binary integer programming.

This approach is particularly simple because you do not give a solution algorithm. Just express the rules of Sudoku, express the clues as constraints on the solution, and then `intlinprog` produces the solution.

Binary Integer Programming Approach

The key idea is to transform a puzzle from a square 9-by-9 grid to a cubic 9-by-9-by-9 array of binary values (0 or 1). Think of the cubic array as being 9 square grids stacked on top of each other. The top grid, a square layer of the array, has a 1 wherever the solution or clue has a 1. The second layer has a 1 wherever the solution or clue has a 2. The ninth layer has a 1 wherever the solution or clue has a 9.

This formulation is precisely suited for binary integer programming.

The objective function is not needed here, and might as well be 0. The problem is really just to find a feasible solution, meaning one that satisfies all the constraints. However, for tiebreaking in the internals of the integer programming solver, giving increased solution speed, use a nonconstant objective function.

Express the Rules for Sudoku as Constraints

Suppose a solution is represented in a 9-by-9-by-9 binary array. What properties does it have? First, each square in the 2-D grid (i,j) has exactly one value, so there is exactly one nonzero element among the 3-D array entries $x_{i,j,k}$. In other words, for every i and j ,

Similarly, in each row i of the 2-D grid, there is exactly one value out of each of the digits from 1 to 9. In other words, for each i and d ,

And each column j in the 2-D grid has the same property: for each j and d ,

The major 3-by-3 grids have a similar constraint. For the grid elements $x_{i,j}$ and $x_{i',j'}$, and for each d ,

To represent all nine major grids, just add 3 or 6 to each i and j index:

where

Express Clues

Each initial value (clue) can be expressed as a constraint. Suppose that the clue is d for some (i,j) . Then $x_{i,j} = d$. The constraint ensures that all other $x_{i',j'} = d$.

Write the Rules for Sudoku

Although the Sudoku rules are conveniently expressed in terms of a 9-by-9-by-9 solution array x , linear constraints are given in terms of a vector solution matrix $x(:)$. Therefore, when you write a Sudoku program, you have to use constraint matrices derived from 9-by-9-by-9 initial arrays.

Here is one approach to set up Sudoku rules, and also include the clues as constraints. The `sudokuEngine` file comes with your software.

```
type sudokuEngine
```

```
function [S,eflag] = sudokuEngine(B)
% This function sets up the rules for Sudoku. It reads in the puzzle
% expressed in matrix B, calls intlinprog to solve the puzzle, and returns
% the solution in matrix S.
```

```

%
% The matrix B should have 3 columns and at least 17 rows (because a Sudoku
% puzzle needs at least 17 entries to be uniquely solvable). The first two
% elements in each row are the i,j coordinates of a clue, and the third
% element is the value of the clue, an integer from 1 to 9. If B is a
% 9-by-9 matrix, the function first converts it to 3-column form.

% Copyright 2014 The MathWorks, Inc.

if isequal(size(B),[9,9]) % 9-by-9 clues
    % Convert to 81-by-3
    [SM,SN] = meshgrid(1:9); % make i,j entries
    B = [SN(:),SM(:),B(:)]; % i,j,k rows
    % Now delete zero rows
    [rrem,~] = find(B(:,3) == 0);
    B(rrem,:) = [];
end

if size(B,2) ~= 3 || length(size(B)) > 2
    error('The input matrix must be N-by-3 or 9-by-9')
end

if sum([any(B ~= round(B)),any(B < 1),any(B > 9)]) % enforces entries 1-9
    error('Entries must be integers from 1 to 9')
end

%% The rules of Sudoku:
N = 9^3; % number of independent variables in x, a 9-by-9-by-9 array
M = 4*9^2; % number of constraints, see the construction of Aeq
Aeq = zeros(M,N); % allocate equality constraint matrix Aeq*x = beq
beq = ones(M,1); % allocate constant vector beq
f = (1:N)'; % the objective can be anything, but having nonconstant f can s
lb = zeros(9,9,9); % an initial zero array
ub = lb+1; % upper bound array to give binary variables

counter = 1;
for j = 1:9 % one in each row
    for k = 1:9
        Astuff = lb; % clear Astuff
        Astuff(1:end,j,k) = 1; % one row in Aeq*x = beq
    end
end

```

```

        Aeq(counter,:) = Astuff(:)'; % put Astuff in a row of Aeq
        counter = counter + 1;
    end
end

for i = 1:9 % one in each column
    for k = 1:9
        Astuff = lb;
        Astuff(i,1:end,k) = 1;
        Aeq(counter,:) = Astuff(:)';
        counter = counter + 1;
    end
end

for U = 0:3:6 % one in each square
    for V = 0:3:6
        for k = 1:9
            Astuff = lb;
            Astuff(U+(1:3),V+(1:3),k) = 1;
            Aeq(counter,:) = Astuff(:)';
            counter = counter + 1;
        end
    end
end

for i = 1:9 % one in each depth
    for j = 1:9
        Astuff = lb;
        Astuff(i,j,1:end) = 1;
        Aeq(counter,:) = Astuff(:)';
        counter = counter + 1;
    end
end

%% Put the particular puzzle in the constraints
% Include the initial clues in the |lb| array by setting corresponding
% entries to 1. This forces the solution to have |x(i,j,k) = 1|.

for i = 1:size(B,1)
    lb(B(i,1),B(i,2),B(i,3)) = 1;
end

```

```

end

%% Solve the Puzzle
% The Sudoku problem is complete: the rules are represented in the |Aeq|
% and |beq| matrices, and the clues are ones in the |lb| array. Solve the
% problem by calling |intlinprog|. Ensure that the integer program has all
% binary variables by setting the intcon argument to |1:N|, with lower and
% upper bounds of 0 and 1.

intcon = 1:N;

[x,~,eflag] = intlinprog(f,intcon,[],[],Aeq,beq,lb,ub);

%% Convert the Solution to a Usable Form
% To go from the solution x to a Sudoku grid, simply add up the numbers at
% each $(i,j)$ entry, multiplied by the depth at which the numbers appear:

if eflag > 0 % good solution
    x = reshape(x,9,9,9); % change back to a 9-by-9-by-9 array
    x = round(x); % clean up non-integer solutions
    y = ones(size(x));
    for k = 2:9
        y(:,:,k) = k; % multiplier for each depth k
    end

    S = x.*y; % multiply each entry by its depth
    S = sum(S,3); % S is 9-by-9 and holds the solved puzzle
else
    S = [];
end

```

Call the Sudoku Solver

```

S = sudokuEngine(B); % Solves the puzzle pictured at the start
drawSudoku(S)

```

LP: Optimal objective value is 29565.000000.

Cut Generation: Applied 2 strong CG cuts,
and 1 zero-half cut.

```
Lower bound is 29565.000000.  
Relative gap is 0.00%.
```

Optimal solution found.

```
Intlinprog stopped at the root node because the objective value is within a  
tolerance of the optimal value; options.TolGapAbs = 0 (the default value).  
intcon variables are integer within tolerance, options.TolInteger = 1e-05 (  
default value).
```

You can easily check that the solution is correct.

Function to Draw the Sudoku Puzzle

```
type drawSudoku
```

```
function drawSudoku(B)  
% Function for drawing the Sudoku board  
  
% Copyright 2014 The MathWorks, Inc.  
  
figure;hold on;axis off;axis equal % prepare to draw  
rectangle('Position',[0 0 9 9],'LineWidth',3,'Clipping','off') % outside bo  
rectangle('Position',[3,0,3,9],'LineWidth',2) % heavy vertical lines  
rectangle('Position',[0,3,9,3],'LineWidth',2) % heavy horizontal lines  
rectangle('Position',[0,1,9,1],'LineWidth',1) % minor horizontal lines  
rectangle('Position',[0,4,9,1],'LineWidth',1)  
rectangle('Position',[0,7,9,1],'LineWidth',1)  
rectangle('Position',[1,0,1,9],'LineWidth',1) % minor vertical lines  
rectangle('Position',[4,0,1,9],'LineWidth',1)  
rectangle('Position',[7,0,1,9],'LineWidth',1)  
  
% Fill in the clues  
%
```



```
% The rows of B are of the form (i,j,k) where i is the row counting from
% the top, j is the column, and k is the clue. To place the entries in the
% boxes, j is the horizontal distance, 10-i is the vertical distance, and
% we subtract 0.5 to center the clue in the box.
%
% If B is a 9-by-9 matrix, convert it to 3 columns first

if size(B,2) == 9 % 9 columns
    [SM,SN] = meshgrid(1:9); % make i,j entries
    B = [SN(:),SM(:),B(:)]; % i,j,k rows
end

for ii = 1:size(B,1)
    text(B(ii,2)-0.5,9.5-B(ii,1),num2str(B(ii,3)))
end

hold off

end
```

Least-Squares (Model Fitting) Algorithms

In this section...

“Least Squares Definition” on page 6-200

“Trust-Region-Reflective Least Squares” on page 6-201

“Levenberg-Marquardt Method” on page 6-205

Least Squares Definition

Least squares, in general, is the problem of finding a vector x that is a local minimizer to a function that is a sum of squares, possibly subject to some constraints:

$$\min_x \|F(x)\|_2^2 = \min_x \sum_i F_i^2(x)$$

such that $Ax \leq b$, $Aeq\ x = beq$, $lb \leq x \leq ub$.

There are several Optimization Toolbox solvers available for various types of $F(x)$ and various types of constraints:

Solver	$F(x)$	Constraints
<code>\</code>	$Cx - d$	None
<code>lsqnonneg</code>	$Cx - d$	$x \geq 0$
<code>lsqlin</code>	$Cx - d$	Bound, linear
<code>lsqnonlin</code>	General $F(x)$	Bound
<code>lsqcurvefit</code>	$F(x, xdata) - ydata$	Bound

There are four least-squares algorithms in Optimization Toolbox solvers, in addition to the algorithms used in `\`:

- Trust-region-reflective
- Levenberg-Marquardt
- `lsqlin` active-set

- The algorithm used by `lsqnonneg`

The trust-region reflective algorithm, `lsqnonneg` algorithm, and Levenberg-Marquardt algorithm are large-scale; see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12. The medium-scale `lsqlin` algorithm is not large-scale. For a general survey of nonlinear least-squares methods, see Dennis [8]. Specific details on the Levenberg-Marquardt method can be found in Moré [28].

Trust-Region-Reflective Least Squares

Trust-Region-Reflective Least Squares Algorithm

Many of the methods used in Optimization Toolbox solvers are based on *trust regions*, a simple yet powerful concept in optimization.

To understand the trust-region approach to optimization, consider the unconstrained minimization problem, minimize $f(x)$, where the function takes vector arguments and returns scalars. Suppose you are at a point x in n -space and you want to improve, i.e., move to a point with a lower function value. The basic idea is to approximate f with a simpler function q , which reasonably reflects the behavior of function f in a neighborhood N around the point x . This neighborhood is the trust region. A trial step s is computed by minimizing (or approximately minimizing) over N . This is the trust-region subproblem,

$$\min_s \{q(s), s \in N\}. \quad (6-99)$$

The current point is updated to be $x + s$ if $f(x + s) < f(x)$; otherwise, the current point remains unchanged and N , the region of trust, is shrunk and the trial step computation is repeated.

The key questions in defining a specific trust-region approach to minimizing $f(x)$ are how to choose and compute the approximation q (defined at the current point x), how to choose and modify the trust region N , and how accurately to solve the trust-region subproblem. This section focuses on the unconstrained problem. Later sections discuss additional complications due to the presence of constraints on the variables.

In the standard trust-region method ([48]), the quadratic approximation q is defined by the first two terms of the Taylor approximation to F at x ; the neighborhood N is usually spherical or ellipsoidal in shape. Mathematically the trust-region subproblem is typically stated

$$\min \left\{ \frac{1}{2} s^T H s + s^T g \text{ such that } \|D s\| \leq \Delta \right\}, \quad (6-100)$$

where g is the gradient of f at the current point x , H is the Hessian matrix (the symmetric matrix of second derivatives), D is a diagonal scaling matrix, Δ is a positive scalar, and $\| \cdot \|$ is the 2-norm. Good algorithms exist for solving Equation 6-100 (see [48]); such algorithms typically involve the computation of a full eigensystem and a Newton process applied to the secular equation

$$\frac{1}{\Delta} - \frac{1}{\|s\|} = 0.$$

Such algorithms provide an accurate solution to Equation 6-100. However, they require time proportional to several factorizations of H . Therefore, for trust-region problems a different approach is needed. Several approximation and heuristic strategies, based on Equation 6-100, have been proposed in the literature ([42] and [50]). The approximation approach followed in Optimization Toolbox solvers is to restrict the trust-region subproblem to a two-dimensional subspace S ([39] and [42]). Once the subspace S has been computed, the work to solve Equation 6-100 is trivial even if full eigenvalue/eigenvector information is needed (since in the subspace, the problem is only two-dimensional). The dominant work has now shifted to the determination of the subspace.

The two-dimensional subspace S is determined with the aid of a preconditioned conjugate gradient process described below. The solver defines S as the linear space spanned by s_1 and s_2 , where s_1 is in the direction of the gradient g , and s_2 is either an approximate Newton direction, i.e., a solution to

$$H \cdot s_2 = -g, \quad (6-101)$$

or a direction of negative curvature,

$$s_2^T \cdot H \cdot s_2 < 0. \quad (6-102)$$

The philosophy behind this choice of S is to force global convergence (via the steepest descent direction or negative curvature direction) and achieve fast local convergence (via the Newton step, when it exists).

A sketch of unconstrained minimization using trust-region ideas is now easy to give:

- 1 Formulate the two-dimensional trust-region subproblem.
- 2 Solve Equation 6-100 to determine the trial step s .
- 3 If $f(x + s) < f(x)$, then $x = x + s$.
- 4 Adjust Δ .

These four steps are repeated until convergence. The trust-region dimension Δ is adjusted according to standard rules. In particular, it is decreased if the trial step is not accepted, i.e., $f(x + s) \geq f(x)$. See [46] and [49] for a discussion of this aspect.

Optimization Toolbox solvers treat a few important special cases of f with specialized functions: nonlinear least-squares, quadratic functions, and linear least-squares. However, the underlying algorithmic ideas are the same as for the general case. These special cases are discussed in later sections.

Large Scale Nonlinear Least Squares

An important special case for $f(x)$ is the nonlinear least-squares problem

$$\min_x \sum_i f_i^2(x) = \min_x \|F(x)\|_2^2, \quad (6-103)$$

where $F(x)$ is a vector-valued function with component i of $F(x)$ equal to $f_i(x)$. The basic method used to solve this problem is the same as in the general case described in “Trust-Region Methods for Nonlinear Minimization” on page 6-5. However, the structure of the nonlinear least-squares problem is exploited to enhance efficiency. In particular, an approximate Gauss-Newton direction, i.e., a solution s to

$$\min \|Js + F\|_2^2, \quad (6-104)$$

(where J is the Jacobian of $F(x)$) is used to help define the two-dimensional subspace S . Second derivatives of the component function $f_i(x)$ are not used.

In each iteration the method of preconditioned conjugate gradients is used to approximately solve the normal equations, i.e.,

$$J^T J s = -J^T F,$$

although the normal equations are not explicitly formed.

Large Scale Linear Least Squares

In this case the function $f(x)$ to be solved is

$$f(x) = \|Cx + d\|_2^2,$$

possibly subject to linear constraints. The algorithm generates strictly feasible iterates converging, in the limit, to a local solution. Each iteration involves the approximate solution of a large linear system (of order n , where n is the length of x). The iteration matrices have the structure of the matrix C . In particular, the method of preconditioned conjugate gradients is used to approximately solve the normal equations, i.e.,

$$C^T C x = -C^T d,$$

although the normal equations are not explicitly formed.

The subspace trust-region method is used to determine a search direction. However, instead of restricting the step to (possibly) one reflection step, as in the nonlinear minimization case, a piecewise reflective line search is conducted at each iteration, as in the quadratic case. See [45] for details of the line search. Ultimately, the linear systems represent a Newton approach capturing the first-order optimality conditions at the solution, resulting in strong local convergence rates.

Jacobian Multiply Function. `lsqlin` can solve the linearly-constrained least-squares problem without using the matrix C explicitly. Instead, it uses a Jacobian multiply function `jmfun`,

```
W = jmfun(Jinfo,Y,flag)
```

that you provide. The function must calculate the following products for a matrix Y :

- If `flag == 0` then $W = C' * (C * Y)$.
- If `flag > 0` then $W = C * Y$.
- If `flag < 0` then $W = C' * Y$.

This can be useful if C is large, but contains enough structure that you can write `jmfun` without forming C explicitly. For an example, see “Jacobian Multiply Function with Linear Least Squares” on page 6-224.

Levenberg-Marquardt Method

In the least-squares problem a function $f(x)$ is minimized that is a sum of squares.

$$\min_x f(x) = \|F(x)\|_2^2 = \sum_i F_i^2(x). \quad (6-105)$$

Problems of this type occur in a large number of practical applications, especially when fitting model functions to data, i.e., nonlinear parameter estimation. They are also prevalent in control where you want the output, $y(x,t)$, to follow some continuous model trajectory, $\varphi(t)$, for vector x and scalar t . This problem can be expressed as

$$\min_{x \in \mathfrak{X}^n} \int_{t_1}^{t_2} (y(x,t) - \varphi(t))^2 dt, \quad (6-106)$$

where $y(x,t)$ and $\varphi(t)$ are scalar functions.

When the integral is discretized using a suitable quadrature formula, the above can be formulated as a least-squares problem:

$$\min_{x \in \mathfrak{R}^n} f(x) = \sum_{i=1}^m (\bar{y}(x, t_i) - \bar{\varphi}(t_i))^2, \quad (6-107)$$

where \bar{y} and $\bar{\varphi}$ include the weights of the quadrature scheme. Note that in this problem the vector $F(x)$ is

$$F(x) = \begin{bmatrix} \bar{y}(x, t_1) - \bar{\varphi}(t_1) \\ \bar{y}(x, t_2) - \bar{\varphi}(t_2) \\ \dots \\ \bar{y}(x, t_m) - \bar{\varphi}(t_m) \end{bmatrix}.$$

In problems of this kind, the residual $\|F(x)\|$ is likely to be small at the optimum since it is general practice to set realistically achievable target trajectories. Although the function in LS can be minimized using a general unconstrained minimization technique, as described in “Basics of Unconstrained Optimization” on page 6-8, certain characteristics of the problem can often be exploited to improve the iterative efficiency of the solution procedure. The gradient and Hessian matrix of LS have a special structure.

Denoting the m -by- n Jacobian matrix of $F(x)$ as $J(x)$, the gradient vector of $f(x)$ as $G(x)$, the Hessian matrix of $f(x)$ as $H(x)$, and the Hessian matrix of each $F_i(x)$ as $H_i(x)$, you have

$$\begin{aligned} G(x) &= 2J(x)^T F(x) \\ H(x) &= 2J(x)^T J(x) + 2Q(x), \end{aligned} \quad (6-108)$$

where

$$Q(x) = \sum_{i=1}^m F_i(x) \cdot H_i(x).$$

The matrix $Q(x)$ has the property that when the residual $\|F(x)\|$ tends to zero as x_k approaches the solution, then $Q(x)$ also tends to zero. Thus when $\|F(x)\|$ is small at the solution, a very effective method is to use the Gauss-Newton direction as a basis for an optimization procedure.

In the Gauss-Newton method, a search direction, d_k , is obtained at each major iteration, k , that is a solution of the linear least-squares problem:

$$\min_{x \in \mathfrak{X}^n} \|J(x_k) - F(x_k)\|_2^2. \quad (6-109)$$

The direction derived from this method is equivalent to the Newton direction when the terms of $Q(x)$ can be ignored. The search direction d_k can be used as part of a line search strategy to ensure that at each iteration the function $f(x)$ decreases.

The Gauss-Newton method often encounters problems when the second-order term $Q(x)$ is significant. A method that overcomes this problem is the Levenberg-Marquardt method.

The Levenberg-Marquardt [25], and [27] method uses a search direction that is a solution of the linear set of equations

$$\left(J(x_k)^T J(x_k) + \lambda_k I \right) d_k = -J(x_k)^T F(x_k), \quad (6-110)$$

or, optionally, of the equations

$$\left(J(x_k)^T J(x_k) + \lambda_k \text{diag} \left(J(x_k)^T J(x_k) \right) \right) d_k = -J(x_k)^T F(x_k), \quad (6-111)$$

where the scalar λ_k controls both the magnitude and direction of d_k . Set option `ScaleProblem` to 'none' to choose Equation 6-110, and set `ScaleProblem` to 'Jacobian' to choose Equation 6-111.

When λ_k is zero, the direction d_k is identical to that of the Gauss-Newton method. As λ_k tends to infinity, d_k tends towards the steepest descent direction, with magnitude tending to zero. This implies that for some sufficiently large λ_k , the term $F(x_k + d_k) < F(x_k)$ holds true. The term λ_k can therefore be controlled to ensure descent even when second-order terms, which restrict the efficiency of the Gauss-Newton method, are encountered.

The Levenberg-Marquardt method therefore uses a search direction that is a cross between the Gauss-Newton direction and the steepest descent direction. This is illustrated in Figure 6-4, Levenberg-Marquardt Method on

Rosenbrock's Function. The solution for Rosenbrock's function converges after 90 function evaluations compared to 48 for the Gauss-Newton method. The poorer efficiency is partly because the Gauss-Newton method is generally more effective when the residual is zero at the solution. However, such information is not always available beforehand, and the increased robustness of the Levenberg-Marquardt method compensates for its occasional poorer efficiency.





Figure 6-4: Levenberg-Marquardt Method on Rosenbrock's Function

For an animated version of this figure, enter `bandem` at the MATLAB command line.

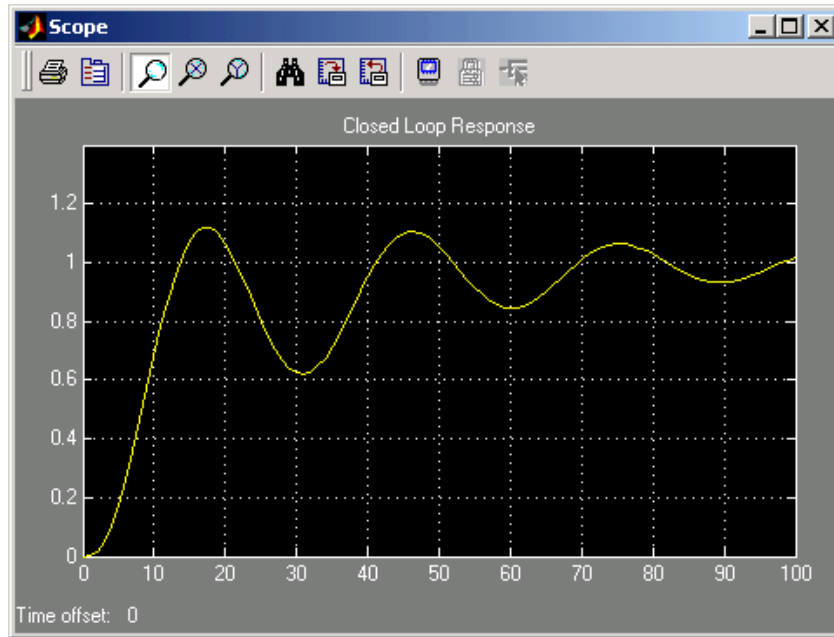
Isqnonlin with a Simulink Model

Suppose that you want to optimize the control parameters in the Simulink model `optsim.mdl`. (This model can be found in the `optim/optimdemos` folder. Note that Simulink must be installed on your system to load this model.) The model includes a nonlinear process plant modeled as a Simulink block diagram.



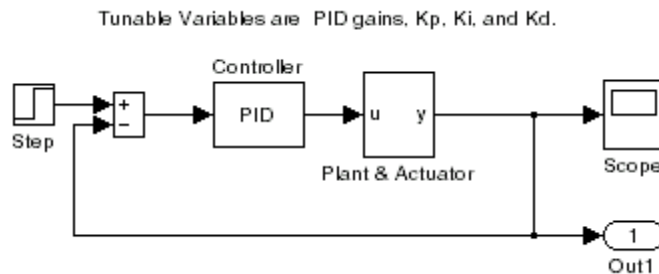
Plant with Actuator Saturation

The plant is an under-damped third-order model with actuator limits. The actuator limits are a saturation limit and a slew rate limit. The actuator saturation limit cuts off input values greater than 2 units or less than -2 units. The slew rate limit of the actuator is 0.8 units/sec. The closed-loop response of the system to a step input is shown in Closed-Loop Response on page 6-210. You can see this response by opening the model (type `optsim` at the command line or click the model name), and selecting **Start** from the **Simulation** menu. The response plots to the scope.



Closed-Loop Response

The problem is to design a feedback control loop that tracks a unit step input to the system. The closed-loop plant is entered in terms of the blocks where the plant and actuator have been placed in a hierarchical Subsystem block. A Scope block displays output trajectories during the design process.



Closed-Loop Model

One way to solve this problem is to minimize the error between the output and the input signal. The variables are the parameters of the Proportional Integral Derivative (PID) controller. If you only need to minimize the error at one time unit, it would be a single objective function. But the goal is to minimize the error for all time steps from 0 to 100, thus producing a multiobjective function (one function for each time step).

The routine `lsqnonlin` is used to perform a least-squares fit on the tracking of the output. The tracking is performed via the function `tracklsq`, which returns the error signal `yout`, the output computed by calling `sim`, minus the input signal `1`. The code for `tracklsq` is contained in the file `runtracklsq.m`, shown below.

The function `runtracklsq` sets up all the needed values and then calls `lsqnonlin` with the objective function `tracklsq`, which is nested inside `runtracklsq`. The variable options passed to `lsqnonlin` defines the criteria and display characteristics. In this case you ask for output, use the medium-scale algorithm, and give termination tolerances for the step and objective function on the order of 0.001.

To run the simulation in the model `optsim`, the variables `Kp`, `Ki`, `Kd`, `a1`, and `a2` (`a1` and `a2` are variables in the Plant block) must all be defined. `Kp`, `Ki`, and `Kd` are the variables to be optimized. The function `tracklsq` is nested inside `runtracklsq` so that the variables `a1` and `a2` are shared between the two functions. The variables `a1` and `a2` are initialized in `runtracklsq`.

The objective function `tracklsq` runs the simulation. The simulation can be run either in the base workspace or the current workspace, that is, the workspace of the function calling `sim`, which in this case is the workspace of `tracklsq`. In this example, the `SrcWorkspace` option is set to 'Current' to tell `sim` to run the simulation in the current workspace. The simulation is performed to 100 seconds.

When the simulation is completed, the `myobj` object is created in the current workspace (that is, the workspace of `tracklsq`). The `Outport` block in the block diagram model puts the `yout` field of the object into the current workspace at the end of the simulation.

The following is the code for `runtracklsq`:

```

function [Kp,Ki,Kd] = runtracklsq
% RUNTRACKLSQ demonstrates using LSQNONLIN with Simulink.

optsim                % Load the model
pid0 = [0.63 0.0504 1.9688]; % Set initial values
a1 = 3; a2 = 43;        % Initialize model plant variables
options = optimoptions(@lsqnonlin,'Algorithm','levenberg-marquardt',...
    'Display','off','TolX',0.001,'TolFun',0.001);
pid = lsqnonlin(@tracklsq, pid0, [], [], options);
Kp = pid(1); Ki = pid(2); Kd = pid(3);

function F = tracklsq(pid)
    % Track the output of optsim to a signal of 1

    % Variables a1 and a2 are needed by the model optsim.
    % They are shared with RUNTRACKLSQ so do not need to be
    % redefined here.
    Kp = pid(1);
    Ki = pid(2);
    Kd = pid(3);

    % Set sim options and compute function value
    myobj = sim('optsim','SrcWorkspace','Current', ...
        'StopTime','100');
    F = myobj.get('yout') - 1;
end
end

```

Copy the code for `runtracklsq` to a file named `runtracklsq.m`, placed in a folder on your MATLAB path.

When you run `runtracklsq`, the optimization gives the solution for the proportional, integral, and derivative (Kp, Ki, Kd) gains of the controller:

```

[Kp, Ki, Kd] = runtracklsq

Done initializing optsim.

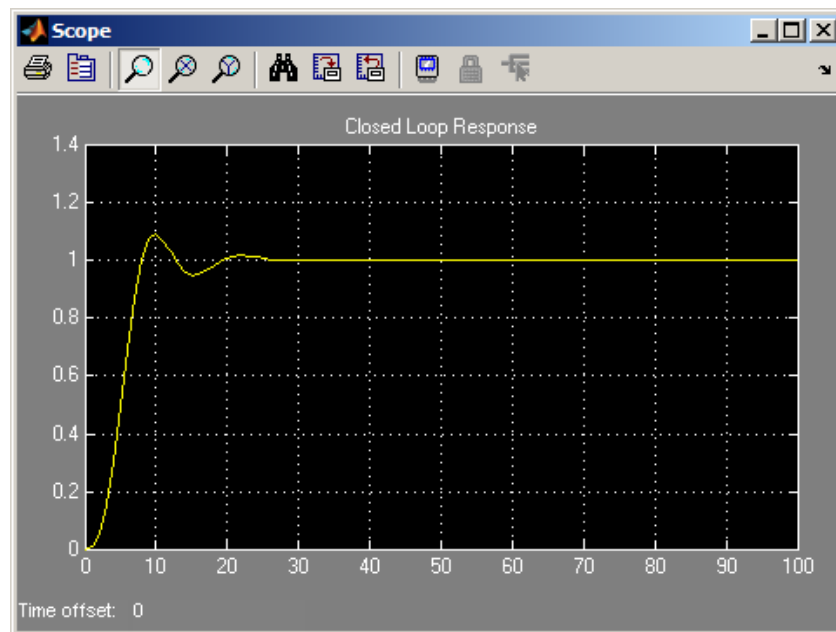
Kp =
    2.9633

```

$K_i =$
0.1436

$K_d =$
13.1386

Here is the resulting closed-loop step response.



Closed-Loop Response Using Isqnonlin

Note The call to `sim` results in a call to one of the Simulink ordinary differential equation (ODE) solvers. A choice must be made about the type of solver to use. From the optimization point of view, a fixed-step solver is the best choice if that is sufficient to solve the ODE. However, in the case of a stiff system, a variable-step method might be required to solve the ODE.

The numerical solution produced by a variable-step solver, however, is not a smooth function of parameters, because of step-size control mechanisms. This lack of smoothness can prevent the optimization routine from converging. The lack of smoothness is not introduced when a fixed-step solver is used. (For a further explanation, see [53].)

Simulink Design Optimization software is recommended for solving multiobjective optimization problems in conjunction with Simulink variable-step solvers. It provides a special numeric gradient computation that works with Simulink and avoids introducing a problem of lack of smoothness.

Nonlinear Least Squares With and Without Jacobian

In this section...

“Jacobians and Memory in Nonlinear Least Squares” on page 6-215

“Step 1: Write a file myfun.m that computes the objective function values.” on page 6-216

“Step 2: Call the nonlinear least-squares routine.” on page 6-216

“Step 3: Include a Jacobian.” on page 6-217

Jacobians and Memory in Nonlinear Least Squares

You can use the trust-region reflective algorithm in `lsqnonlin`, `lsqcurvefit`, and `fsolve` with small- to medium-scale problems without computing the Jacobian in `fun` or providing the Jacobian sparsity pattern. (This example also applies when using `fmincon` or `fminunc` without computing the Hessian or supplying the Hessian sparsity pattern.) How small is small- to medium-scale? No absolute answer is available, as it depends on the amount of virtual memory in your computer system configuration.

Suppose your problem has m equations and n unknowns. If the command `J = sparse(ones(m,n))` causes an `Out of memory` error on your machine, then this is certainly too large a problem. If it does not result in an error, the problem might still be too large. You can only find out by running it and seeing if MATLAB runs within the amount of virtual memory available on your system.

Suppose you have a small problem with 10 equations and 2 unknowns, such as finding x that minimizes

$$\sum_{k=1}^{10} \left(2 + 2k - e^{kx_1} - e^{kx_2} \right)^2,$$

starting at the point $x = [0.3, 0.4]$.

Because `lsqnonlin` assumes that the sum of squares is not explicitly formed in the user function, the function passed to `lsqnonlin` should instead compute the vector valued function

$$F_k(x) = 2 + 2k - e^{kx_1} - e^{kx_2},$$

for $k = 1$ to 10 (that is, F should have 10 components).

Step 1: Write a file `myfun.m` that computes the objective function values.

```
function F = myfun(x)
k = 1:10;
F = 2 + 2*k - exp(k*x(1)) - exp(k*x(2));
```

Step 2: Call the nonlinear least-squares routine.

```
x0 = [0.3 0.4];           % Starting guess
[x,resnorm,res,eflag,output1] = lsqnonlin(@myfun,x0);
% Invoke optimizer
```

Because the Jacobian is not computed in `myfun.m`, and no Jacobian sparsity pattern is provided by the `JacobPattern` option in `options`, `lsqnonlin` calls the trust-region reflective algorithm with `JacobPattern` set to `Jstr = sparse(ones(10,2))`. This is the default for `lsqnonlin`. Note that the Jacobian option in `options` is set to `'off'` by default.

When the finite-differencing routine is called initially, it detects that `Jstr` is actually a dense matrix, i.e., no speed benefit is derived from storing it as a sparse matrix. From then on, the finite-differencing routine uses `Jstr = ones(10,2)` (a full matrix) for the optimization computations.

After about 24 function evaluations, this example gives the solution

```
x,resnorm
x =
    0.2578    0.2578
```

```
resnorm =
    124.3622
```

Most computer systems can handle much larger full problems, say into the hundreds of equations and variables. But *if* there is some sparsity structure in the Jacobian (or Hessian) that can be taken advantage of, the large-scale methods always runs faster if this information is provided.

Step 3: Include a Jacobian.

The objective function is simple enough to calculate its Jacobian. Following the definition in “Jacobians of Vector Functions” on page 2-27, a Jacobian function represents the matrix

$$J_{kj}(x) = \frac{\partial F_k(x)}{\partial x_j}.$$

Here, $F_k(x)$ is the k th component of the objective function. This example has

$$F_k(x) = 2 + 2k - e^{kx_1} - e^{kx_2},$$

so

$$J_{k1}(x) = -ke^{kx_1}$$

$$J_{k2}(x) = -ke^{kx_2}.$$

Modify the objective function file.

```
function [F,J] = myfun(x)
k = 1:10;
F = 2 + 2*k - exp(k*x(1)) - exp(k*x(2));
if nargin > 1
    J = zeros(10,2);
    J(k,1) = -k.*exp(k*x(1));
    J(k,2) = -k.*exp(k*x(2));
end
```

Set options so the solver uses the Jacobian.

```
opts = optimoptions(@lsqnonlin,'Jacobian','on');
```

Run the solver.

```
x0 = [0.3 0.4];           % Starting guess  
[x,resnorm,res,eflag,output2] = lsqnonlin(@myfun,x0,[],[],opts);
```

The solution is the same as before.

```
x,resnorm
```

```
x =  
    0.2578    0.2578
```

```
resnorm =  
    124.3622
```

The advantage to using a Jacobian is that the solver takes fewer function evaluations, 24 instead of 72.

```
[output1.funcCount,output2.funcCount]
```

```
ans =  
    72    24
```

Linear Least Squares with Bound Constraints

Many situations give rise to sparse linear least-squares problems, often with bounds on the variables. The next problem requires that the variables be nonnegative. This problem comes from fitting a function approximation to a piecewise linear spline. Specifically, particles are scattered on the unit square. The function to be approximated is evaluated at these points, and a piecewise linear spline approximation is constructed under the condition that (linear) coefficients are not negative. There are 2000 equations to fit on 400 variables:

```
load particle % Get C, d
lb = zeros(400,1);
[x,resnorm,residual,exitflag,output] = ...
    lsqlin(C,d,[],[],[],[],lb);
```

The default diagonal preconditioning works fairly well:

```
exitflag,resnorm,output
```

```
exitflag =
    3
```

```
resnorm =
    22.5794
```

```
output =
    iterations: 10
    algorithm: 'large-scale: trust-region reflective Newton'
    firstorderopt: 2.7870e-005
    cgiterations: 42
    message: 'Optimization terminated: relative function value change
    than sqr...'
```

For bound constrained problems, the first-order optimality is the infinity norm of $v \cdot *g$, where v is defined as in “Box Constraints” on page 6-30, and g is the gradient.

You can improve (decrease) the first-order optimality measure by using a sparse QR factorization in each iteration. To do this, set `PrecondBandWidth` to `inf`:

```
options = optimoptions('lsqlin','PrecondBandWidth',inf);
[x,resnorm,residual,exitflag,output] = ...
    lsqlin(C,d,[],[],[],[],lb,[],[],options);
```

The first-order optimality measure decreases:

```
exitflag,resnorm,output
```

```
exitflag =
    1
```

```
resnorm =
    22.5794
```

```
output =
    iterations: 12
    algorithm: 'large-scale: trust-region reflective Newton'
    firstorderopt: 5.5907e-015
    cgiterations: 0
    message: 'Optimization terminated: first order optimality with op
gradient n...'
```

Optimization App with the lsqlin Solver

In this section...

“The Problem” on page 6-221

“Setting Up the Problem” on page 6-221

The Problem

This example shows how to use the Optimization app to solve a constrained least-squares problem.

The problem in this example is to find the point on the plane $x_1 + 2x_2 + 4x_3 = 7$ that is closest to the origin. The easiest way to solve this problem is to minimize the square of the distance from a point $x = (x_1, x_2, x_3)$ on the plane to the origin, which returns the same optimal point as minimizing the actual distance. Since the square of the distance from an arbitrary point (x_1, x_2, x_3) to the origin is $x_1^2 + x_2^2 + x_3^2$, you can describe the problem as follows:

$$\min_x f(x) = x_1^2 + x_2^2 + x_3^2,$$

subject to the constraint

$$x_1 + 2x_2 + 4x_3 = 7.$$

The function $f(x)$ is called the *objective function* and $x_1 + 2x_2 + 4x_3 = 7$ is an *equality constraint*. More complicated problems might contain other equality constraints, inequality constraints, and upper or lower bound constraints.

Setting Up the Problem

This section shows how to set up the problem with the lsqlin solver in the Optimization app.

- 1 Enter `optimtool` in the Command Window to open the Optimization app.
- 2 Select `lsqlin` from the selection of solvers. Use the default algorithm.

Solver:

Algorithm:

3 Enter the following to create variables for the objective function:

- In the **C** field, enter `eye(3)`.
- In the **d** field, enter `zeros(3,1)`.

The **C** and **d** fields should appear as shown in the following figure.

Problem

C:

d:

4 Enter the following to create variables for the equality constraints:

- In the **Aeq** field, enter `[1 2 4]`.
- In the **beq** field, enter 7.

The **Aeq** and **beq** fields should appear as shown in the following figure.

Constraints:

Linear inequalities: A: b:

Linear equalities: Aeq: beq:

Bounds: Lower: Upper:

5 Click the **Start** button as shown in the following figure.

Run solver and view results

Current iteration:

- 6 When the algorithm terminates, under **Run solver and view results** the following information is displayed:

Run solver and view results

Start Pause Stop

Current iteration: 1 Clear Results

Optimization running.
Objective function value: 2.333333333333333
Optimization terminated.

Final point:

Index ▲	Value
1	0.333
2	0.667
3	1.333

- The **Current iteration** value when the algorithm terminated, which for this example is 1.
- The final value of the objective function when the algorithm terminated:
Objective function value: 2.333333333333333
- The algorithm termination message:
Optimization terminated.
- The final point, which for this example is

0.3333
0.6667
1.3333

Jacobian Multiply Function with Linear Least Squares

You can solve a least-squares problem of the form

$$\min_x \frac{1}{2} \|C \cdot x - d\|_2^2$$

such that $Ax \leq b$, $Aeq\ x = beq$, $lb \leq x \leq ub$, for problems where C is very large, perhaps too large to be stored, by using a Jacobian multiply function.

For example, consider the case where C is a $2n$ -by- n matrix based on a circulant matrix. This means the rows of C are shifts of a row vector v . This example has the row vector v with elements of the form $(-1)^{k+1}/k$:

$$v = [1, -1/2, 1/3, -1/4, \dots, -1/n],$$

cyclically shifted:

$$C = \begin{bmatrix} 1 & -1/2 & 1/3 & \dots & -1/n \\ -1/n & 1 & -1/2 & \dots & 1/(n-1) \\ 1/(n-1) & -1/n & 1 & \dots & -1/(n-2) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1/2 & 1/3 & -1/4 & \dots & 1 \\ 1 & -1/2 & 1/3 & \dots & -1/n \\ -1/n & 1 & -1/2 & \dots & 1/(n-1) \\ 1/(n-1) & -1/n & 1 & \dots & -1/(n-2) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1/2 & 1/3 & -1/4 & \dots & 1 \end{bmatrix}.$$

This least-squares example considers the problem where

$$d = [n-1; n-2; \dots; -n],$$

and the constraints are $-5 \leq x(i) \leq 5$ for $i = 1, \dots, n$.

For large enough n , the dense matrix C does not fit into computer memory. ($n = 10,000$ is too large on one tested system.)

A Jacobian multiply function has the following syntax:

```
w = jmfcn(Jinfo,Y,flag)
```

$Jinfo$ is a matrix the same size as C , used as a preconditioner. If C is too large to fit into memory, $Jinfo$ should be sparse. Y is a vector or matrix sized so that $C*Y$ or $C'*Y$ makes sense. $flag$ tells `jmfcn` which product to form:

- $flag > 0 \Rightarrow w = C*Y$
- $flag < 0 \Rightarrow w = C'*Y$
- $flag = 0 \Rightarrow w = C'*C*Y$

Since C is such a simply structured matrix, it is easy to write a Jacobian multiply function in terms of the vector v ; i.e., without forming C . Each row of $C*Y$ is the product of a shifted version of v times Y . The following matrix performs one step of the shift: v shifts to $v*T$, where

$$T = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ 1 & 0 & 0 & \dots & 0 \end{bmatrix}.$$

To compute $C*Y$, compute $v*Y$ to find the first row, then shift v and compute the second row, and so on.

To compute $C'*Y$, perform the same computation, but use a shifted version of $temp$, the vector formed from the first row of C' :

```
temp = [fliplr(v)*T,fliplr(v)*T];
```

To compute $C'*C*Y$, simply compute $C*Y$ using shifts of v , and then compute C' times the result using shifts of `fliplr(v)`.

The `dolsqJac` function in the following code sets up the vector v and matrix T , and calls the solver `lsqlin`:

```
function [x,resnorm,residual,exitflag,output] = dolsqJac(n)
%
r = 1:n-1; % index for making vectors

T = spalloc(n,n,n); % making a sparse circulant matrix
for m = r
    T(m,m+1)=1;
end
T(n,1) = 1;

v(n) = (-1)^(n+1)/n; % allocating the vector v
v(r) = (-1)^(r+1)./r;

% Now C should be a 2n-by-n circulant matrix based on v,
% but that might be too large to fit into memory.

r = 1:2*n;
d(r) = n-r;

Jinfo = [speye(n);speye(n)]; % sparse matrix for preconditioning
% This matrix is a required input for the solver;
% preconditioning is not really being used in this example

% Pass the matrix T and vector v so they don't need to be
% computed in the Jacobian multiply function
options = optimoptions('lsqlin','JacobMult',...
    @(Jinfo,Y,flag)lsqcirculant(Jinfo,Y,flag,T,v));

lb = -5*ones(1,n);
ub = 5*ones(1,n);

[x,resnorm,residual,exitflag,output] = ...
    lsqlin(Jinfo,d,[],[],[],[],lb,ub,[],options);
```

The Jacobian multiply function `lsqcirculant` is as follows:

```
function w = lsqcirculant(Jinfo,Y,flag,T,v)
% This function computes the Jacobian multiply functions
```

```

% for a 2n-by-n circulant matrix example

if flag > 0
    w = Jpositive(Y);
elseif flag < 0
    w = Jnegative(Y);
else
    w = Jnegative(Jpositive(Y));
end

function a = Jpositive(q)
    % Calculate C*q
    temp = v;

    a = zeros(size(q)); % allocating the matrix a
    a = [a;a]; % the result is twice as tall as the input

    for r = 1:size(a,1)
        a(r,:) = temp*q; % compute the rth row
        temp = temp*T; % shift the circulant
    end
end

function a = Jnegative(q)
    % Calculate C'*q
    temp = fliplr(v)*T; % the circulant for C'

    len = size(q,1)/2; % the returned vector is half as long
    % as the input vector
    a = zeros(len,size(q,2)); % allocating the matrix a

    for r = 1:len
        a(r,:) = [temp,temp]*q; % compute the rth row
        temp = temp*T; % shift the circulant
    end
end
end
end

```

When $n = 3000$, C is an 18,000,000-element dense matrix. Here are the results of the `dolsqJac` function for $n = 3000$ at selected values of x , and the output structure:

```
[x,resnorm,residual,exitflag,output] = dolsqJac(3000);
```

```
Optimization terminated: relative function value changing by  
less than OPTIONS.TolFun.
```

```
x(1)  
ans =  
    5.0000
```

```
x(1500)  
ans =  
   -0.5201
```

```
x(3000)  
ans =  
   -5.0000
```

```
output  
output =  
    iterations: 16  
    algorithm: 'large-scale: trust-region reflective Newton'  
 firstorderopt: 7.5143e-05  
    cgiterations: 36  
    message: 'Optimization terminated: relative function value changi  
than OPT...'
```

Nonlinear Curve Fitting with lsqcurvefit

lsqcurvefit enables you to fit parameterized nonlinear functions to data easily. You can use lsqnonlin as well; lsqcurvefit is simply a convenient way to call lsqnonlin for curve fitting.

In this example, the vector `xdata` represents 100 data points, and the vector `ydata` represents the associated measurements. Generate the data using the following script:

```
rng(5489,'twister') % reproducible
xdata = -2*log(rand(100,1));
ydata = (ones(100,1) + .1*randn(100,1)) + (3*ones(100,1)+...
    0.5*randn(100,1)).*exp(-(2*ones(100,1)+...
    .5*randn(100,1)).*xdata);
```

The modeled relationship between `xdata` and `ydata` is

$$ydata_i = a_1 + a_2 \exp(-a_3 xdata_i) + \varepsilon_i. \quad (6-112)$$

The script generates `xdata` from 100 independent samples from an exponential distribution with mean 2. It generates `ydata` from Equation 6-112 using $\mathbf{a} = [1; 3; 2]$, perturbed by adding normal deviates with standard deviations $[0.1; 0.5; 0.5]$.

The goal is to find parameters \hat{a}_i , $i = 1, 2, 3$, for the model that best fit the data.

In order to fit the parameters to the data using `lsqcurvefit`, you need to define a fitting function. Define the fitting function predicted as an anonymous function:

```
predicted = @(a,xdata) a(1)*ones(100,1)+a(2)*exp(-a(3)*xdata);
```

To fit the model to the data, `lsqcurvefit` needs an initial estimate `a0` of the parameters. Enter

```
a0 = [2;2;2];
```

Run the solver `lsqcurvefit` as follows:

```
[ahat,resnorm,residual,exitflag,output,lambda,jacobian] =...  
    lsqcurvefit(predicted,a0,xdata,ydata);
```

Local minimum possible.

lsqcurvefit stopped because the final change in the sum of squares relative to its initial value is less than the default value of the function tolerance.

To see the resulting least-squares estimate of \hat{a} , enter:

```
ahat  
  
ahat =  
    1.0169  
    3.1444  
    2.1596
```

The fitted values ahat are within 8% of $a = [1;3;2]$.

If you have Statistics Toolbox™ software, use the nlparci function to generate confidence intervals for the ahat estimate.

Fit a Model to Complex-Valued Data

This example shows how to perform nonlinear fitting of complex-valued data. While most Optimization Toolbox solvers and algorithms operate only on real-valued data, the `levenberg-marquardt` algorithm works on both real-valued and complex-valued data.

Do not set the `FunValCheck` option to 'on' when using complex data. The solver errors.

Data Model

The data model is a simple exponential:

$$y(x) = v_1 + v_2 e^{v_3 x}.$$

The x is input data, y is the response, and v is a complex-valued vector of coefficients. The goal is to estimate v from x and noisy observations y .

Artificial Data with Noise

Generate artificial data for the model. Take the complex coefficient vector v as `[2;3+4i;-.5+.4i]`. Take the observations x as exponentially distributed. Add complex-valued noise to the responses y .

```
rng default % for reproducibility
N = 100; % number of observations
v0 = [2;3+4i;-.5+.4i]; % coefficient vector
xdata = -log(rand(N,1)); % exponentially distributed
noisedata = randn(N,1).*exp((1i*randn(N,1))); % complex noise
cplxdata = v0(1) + v0(2).*exp(v0(3)*xdata) + noisedata;
```

Fit the Model to Recover the Coefficient Vector

The difference between the response predicted by the data model and an observation (`xdata` for x and response `cplxdata` for y) is:

```
objfcn = @(v)v(1)+v(2)*exp(v(3)*xdata) - cplxdata;
```

Use either `lsqnonlin` or `lsqcurvefit` to fit the model to the data. This example first uses `lsqnonlin`. Because the data is complex, set the `Algorithm` option to `'levenberg-marquardt'`.

```
opts = optimoptions(@lsqnonlin,...
    'Algorithm','levenberg-marquardt','Display','off');
x0 = (1+1i)*[1;1;1]; % arbitrary initial guess
[vestimated,resnorm,residuals,exitflag,output] = lsqnonlin(objfcn,x0,[],[],
vestimated,resnorm,exitflag,output.firstorderopt
```

```
vestimated =
```

```
    2.1581 + 0.1351i
    2.7399 + 3.8012i
   -0.5338 + 0.4660i
```

```
resnorm =
```

```
    100.9933
```

```
exitflag =
```

```
    3
```

```
ans =
```

```
    0.0013
```

`lsqnonlin` recovers the complex coefficient vector to about one significant digit. The norm of the residual is sizable, indicating that the noise keeps the model from fitting all the observations. The exit flag is 3, not the preferable 1, because the first-order optimality measure is about $1e-3$, not below $1e-6$.

Alternative: Use `lsqcurvefit`

To fit using `lsqcurvefit`, write the model to give just the responses, not the responses minus the response data.

```
objfcn = @(v,xdata)v(1)+v(2)*exp(v(3)*xdata);
```

Use `lsqcurvefit` options and syntax.

```
opts = optimoptions(@lsqcurvefit,opts); % reuse the options
[vestimated,resnorm] = lsqcurvefit(objfcn,x0,xdata,cplxdata,[],[],opts)
```

```
vestimated =

    2.1581 + 0.1351i
    2.7399 + 3.8012i
   -0.5338 + 0.4660i
```

```
resnorm =

    100.9933
```

The results match those from `lsqnonlin`, because the underlying algorithms are identical. Use whichever solver you find more convenient.

Alternative: Split Real and Imaginary Parts

To use the trust-region-reflective algorithm, such as when you want to include bounds, you must split the real and complex parts of the coefficients into separate variables. For this problem, split the coefficients as follows:

$$\begin{aligned} y &= v_1 + iv_2 + (v_3 + iv_4) \exp((v_5 + iv_6)x) \\ &= (v_1 + v_3 \exp(v_5x) \cos(v_6x) - v_4 \exp(v_5x) \sin(v_6x)) \\ &\quad + i(v_2 + v_4 \exp(v_5x) \cos(v_6x) + v_3 \exp(v_5x) \sin(v_6x)). \end{aligned}$$

Write the response function for `lsqcurvefit`.

```
function yout = cplxreal(v,xdata)

yout = zeros(length(xdata),2); % allocate yout
expcoef = exp(v(5)*xdata(:)); % magnitude
coscoef = cos(v(6)*xdata(:)); % real cosine term
sincoef = sin(v(6)*xdata(:)); % imaginary sin term
yout(:,1) = v(1) + expcoef.*(v(3)*coscoef - v(4)*sincoef);
```

```
yout(:,2) = v(2) + expcoef.*(v(4)*coscoef + v(3)*sincoef);
```

Save this code as the file `cplxreal.m` on your MATLAB path.

Split the response data into its real and imaginary parts.

```
ydata2 = [real(cplxdata), imag(cplxdata)];
```

The coefficient vector `v` now has six dimensions. Initialize it as all ones, and solve the problem using `lsqcurvefit`.

```
x0 = ones(6,1);  
[vestimated, resnorm, residuals, exitflag, output] = ...  
    lsqcurvefit(@cplxreal, x0, xdata, ydata2);  
vestimated, resnorm, exitflag, output.firstorderopt
```

```
vestimated =
```

```
    2.1582  
    0.1351  
    2.7399  
    3.8012  
   -0.5338  
    0.4660
```

```
resnorm =
```

```
    100.9933
```

```
exitflag =
```

```
    3
```

```
ans =
```

```
    0.0018
```

Interpret the six-element vector `vestimated` as a three-element complex vector, and you see that the solution is virtually the same as the previous solutions.

Multiobjective Optimization Algorithms

In this section...

“Multiobjective Optimization Definition” on page 6-236

“Algorithms” on page 6-237

Multiobjective Optimization Definition

There are two Optimization Toolbox multiobjective solvers: `fgoalattain` and `fminimax`.

- `fgoalattain` addresses the problem of reducing a set of nonlinear functions $F_i(x)$ below a set of goals F_i^* . Since there are several functions $F_i(x)$, it is not always clear what it means to solve this problem, especially when you cannot achieve all the goals simultaneously. Therefore, the problem is reformulated to one that is always well-defined.

The *unscaled goal attainment problem* is to minimize the maximum of $F_i(x) - F_i^*$.

There is a useful generalization of the unscaled problem. Given a set of positive weights w_i , the *goal attainment problem* tries to find x to minimize the maximum of

$$\frac{F_i(x) - F_i^*}{w_i}. \quad (6-113)$$

This minimization is supposed to be accomplished while satisfying all types of constraints: $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, and $l \leq x \leq u$.

If you set all weights equal to 1 (or any other positive constant), the goal attainment problem is the same as the unscaled goal attainment problem. If the F_i^* are positive, and you set all weights as $w_i = F_i^*$, the goal attainment problem becomes minimizing the relative difference between the functions $F_i(x)$ and the goals F_i^* .

In other words, the goal attainment problem is to minimize a slack variable γ , defined as the maximum over i of the expressions in Equation 6-113.

This implies the expression that is the formal statement of the goal attainment problem:

$$\min_{x, \gamma}$$

such that $F(x) - w\gamma \leq F^*$, $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, and $l \leq x \leq u$.

- `fminimax` addresses the problem of minimizing the maximum of a set of nonlinear functions, subject to all types of constraints:

$$\min_x \max_i F_i(x)$$

such that $c(x) \leq 0$, $ceq(x) = 0$, $Ax \leq b$, $Aeqx = beq$, and $l \leq x \leq u$.

Clearly, this problem is a special case of the unscaled goal attainment problem, with $F_i^* = 0$ and $w_i = 1$.

Algorithms

Goal Attainment Method

This section describes the goal attainment method of Gembicki [16]. This

method uses a set of design goals, $F^* = \{F_1^*, F_2^*, \dots, F_m^*\}$, associated with a set of objectives, $F(x) = \{F_1(x), F_2(x), \dots, F_m(x)\}$. The problem formulation allows the objectives to be under- or overachieved, enabling the designer to be relatively imprecise about the initial design goals. The relative degree of under- or overachievement of the goals is controlled by a vector of weighting coefficients, $w = \{w_1, w_2, \dots, w_m\}$, and is expressed as a standard optimization problem using the formulation

$$\underset{\gamma \in \mathfrak{R}, x \in \Omega}{\text{minimize}} \gamma \tag{6-114}$$

such that $F_i(x) - w_i\gamma \leq F_i^*$, $i = 1, \dots, m$.

The term $w_i\gamma$ introduces an element of *slackness* into the problem, which otherwise imposes that the goals be rigidly met. The weighting vector, w , enables the designer to express a measure of the relative tradeoffs between the objectives. For instance, setting the weighting vector w equal to the initial goals indicates that the same percentage under- or overachievement of the goals, F^* , is achieved. You can incorporate hard constraints into the

design by setting a particular weighting factor to zero (i.e., $w_i = 0$). The goal attainment method provides a convenient intuitive interpretation of the design problem, which is solvable using standard optimization procedures. Illustrative examples of the use of the goal attainment method in control system design can be found in Fleming ([10] and [11]).

The goal attainment method is represented geometrically in the figure below in two dimensions.

$$\begin{aligned} \text{minimize } \gamma \quad & \text{subject to } F_1(x) - w_1\gamma \leq F_1^* \\ \gamma, x \in \Omega \quad & F_2(x) - w_2\gamma \leq F_2^* \end{aligned}$$




Figure 6-5: Geometrical Representation of the Goal Attainment Method

Specification of the goals, $\{F_1^*, F_2^*\}$, defines the goal point, P . The weighting vector defines the direction of search from P to the feasible function space, $\Lambda(\gamma)$. During the optimization γ is varied, which changes the size of the feasible region. The constraint boundaries converge to the unique solution point F_{1s}, F_{2s} .

Algorithm Improvements for the Goal Attainment Method

The goal attainment method has the advantage that it can be posed as a nonlinear programming problem. Characteristics of the problem can also be exploited in a nonlinear programming algorithm. In sequential quadratic programming (SQP), the choice of merit function for the line search is not easy because, in many cases, it is difficult to “define” the relative importance between improving the objective function and reducing constraint violations.

This has resulted in a number of different schemes for constructing the merit function (see, for example, Schittkowski [36]). In goal attainment programming there might be a more appropriate merit function, which you can achieve by posing Equation 6-114 as the minimax problem

$$\underset{x \in \mathcal{R}^n}{\text{minimize}} \quad \max_i \{\Lambda_i\}, \quad (6-115)$$

where

$$\Lambda_i = \frac{F_i(x) - F_i^*}{w_i}, \quad i = 1, \dots, m.$$

Following the argument of Brayton et al. [2] for minimax optimization using SQP, using the merit function of Equation 6-47 for the goal attainment problem of Equation 6-115 gives

$$\psi(x, \gamma) = \gamma + \sum_{i=1}^m r_i \cdot \max\{0, F_i(x) - w_i \gamma - F_i^*\}. \quad (6-116)$$

When the merit function of Equation 6-116 is used as the basis of a line search procedure, then, although $\psi(x, \gamma)$ might decrease for a step in a given search direction, the function $\max \Lambda_i$ might paradoxically increase. This is accepting a degradation in the worst case objective. Since the worst case objective is responsible for the value of the objective function γ , this is accepting a step that ultimately increases the objective function to be minimized. Conversely, $\psi(x, \gamma)$ might increase when $\max \Lambda_i$ decreases, implying a rejection of a step that improves the worst case objective.

Following the lines of Brayton et al. [2], a solution is therefore to set $\psi(x)$ equal to the worst case objective, i.e.,

$$\psi(x) = \max_i \Lambda_i. \quad (6-117)$$

A problem in the goal attainment method is that it is common to use a weighting coefficient equal to 0 to incorporate hard constraints. The merit function of Equation 6-117 then becomes infinite for arbitrary violations of the constraints.

To overcome this problem while still retaining the features of Equation 6-117, the merit function is combined with that of Equation 6-48, giving the following:

$$\psi(x) = \sum_{i=1}^m \begin{cases} r_i \cdot \max\{0, F_i(x) - w_i\gamma - F_i^*\} & \text{if } w_i = 0 \\ \max_i \Lambda_i, \quad i = 1, \dots, m & \text{otherwise.} \end{cases} \quad (6-118)$$

Another feature that can be exploited in SQP is the objective function γ . From the KKT equations it can be shown that the approximation to the Hessian of the Lagrangian, H , should have zeros in the rows and columns associated with the variable γ . However, this property does not appear if H is initialized as the identity matrix. H is therefore initialized and maintained to have zeros in the rows and columns associated with γ .

These changes make the Hessian, H , indefinite. Therefore H is set to have zeros in the rows and columns associated with γ , except for the diagonal element, which is set to a small positive number (e.g., 1e-10). This allows use of the fast converging positive definite QP method described in “Quadratic Programming Solution” on page 6-37.

The preceding modifications have been implemented in `fgoalattain` and have been found to make the method more robust. However, because of the rapid convergence of the SQP method, the requirement that the merit function strictly decrease sometimes requires more function evaluations than an implementation of SQP using the merit function of Equation 6-47.

Minimizing the Maximum Objective

`fminimax` uses a goal attainment method. It takes goals of 0, and weights of 1. With this formulation, the goal attainment problem becomes

$$\min_i \max_x \left(\frac{f_i(x) - goal_i}{weight_i} \right) = \min_i \max_x f_i(x),$$

which is the minimax problem.

Parenthetically, you might expect `fminimax` to turn the multiobjective function into a single objective. The function

$$f(x) = \max(F_1(x), \dots, F_j(x))$$

is a single objective function to minimize. However, it is not differentiable, and Optimization Toolbox objectives are required to be smooth. Therefore the minimax problem is formulated as a smooth goal attainment problem.

Using `fminimax` with a Simulink Model

Another approach to optimizing the control parameters in the Simulink model shown in Plant with Actuator Saturation on page 6-209 is to use the `fminimax` function. In this case, rather than minimizing the error between the output and the input signal, you minimize the maximum value of the output at any time t between 0 and 100.

The code for this example, shown below, is contained in the function `runtrackmm`, in which the objective function is simply the output `yout` returned by the `sim` command. But minimizing the maximum output at all time steps might force the output to be far below unity for some time steps. To keep the output above 0.95 after the first 20 seconds, the constraint function `trackmmcon` contains the constraint `yout >= 0.95` from $t=20$ to $t=100$. Because constraints must be in the form $g \leq 0$, the constraint in the function is `g = -yout(20:100)+.95`.

Both `trackmmobj` and `trackmmcon` use the result `yout` from `sim`, calculated from the current PID values. To avoid calling the simulation twice, `runtrackmm` has nested functions so that the value of `yout` is shared between the objective and constraint functions. The simulation is called only when the current point changes.

The following is the code for `runtrackmm`:

```
function [Kp, Ki, Kd] = runtrackmm

    optsim % initialize Simulink(R)
    pid0 = [0.63 0.0504 1.9688];
    % a1, a2, yout are shared with TRACKMMOBJ and TRACKMMCON
    a1 = 3; a2 = 43; % Initialize plant variables in model
    yout = []; % Give yout an initial value
    pold = []; % tracks last pid
    options = optimoptions('fminimax','Display','iter',...
        'TolX',0.001,'TolFun',0.001);
    pid = fminimax(@trackmmobj,pid0,[],[],[],[],[],[],...
        @trackmmcon,options);
    Kp = pid(1); Ki = pid(2); Kd = pid(3);

    function F = trackmmobj(pid)
```

```

        % Track the output of optsim to a signal of 1.
        % Variables a1 and a2 are shared with RUNTRACKMM.
        % Variable yout is shared with RUNTRACKMM and
        % RUNTRACKMMCON.
        updateIfNeeded(pid)

        F = yout;
    end

function [c,ceq] = trackmmcon(pid)
    % Track the output of optsim to a signal of 1.
    % Variable yout is shared with RUNTRACKMM and
    % TRACKMMOBJ
    updateIfNeeded(pid)

    c = -yout(20:100)+.95;
    ceq=[];
end

function updateIfNeeded(pid)
    if ~isequal(pid,pold) % compute only if needed

        Kp = pid(1);
        Ki = pid(2);
        Kd = pid(3);

        myobj = sim('optsim','SrcWorkspace','Current');
        yout = myobj.get('yout');

        pold = pid;
    end
end

end

```

Copy the code for runtrackmm to a file named runtrackmm.m, placed in a folder on your MATLAB path.

When you run the code, it returns the following results:

```
[Kp,Ki,Kd] = runtrackmm
Done initializing optsim.
```

Iter	F-count	Objective value	Max constraint	Line search steplength	Directional derivative	Procedure
0	5	0	1.11982			
1	11	1.184	0.07978	1	0.482	
2	17	1.012	0.04285	1	-0.236	
3	23	0.9996	0.00397	1	-0.0195	Hessian modified twice
4	29	0.9996	3.464e-005	1	0.000687	Hessian modified
5	35	0.9996	2.272e-009	1	-0.0175	Hessian modified twice

Local minimum possible. Constraints satisfied.

fminimax stopped because the size of the current search direction is less than twice the selected value of the step size tolerance and constraints are satisfied to within the default value of the constraint tolerance.

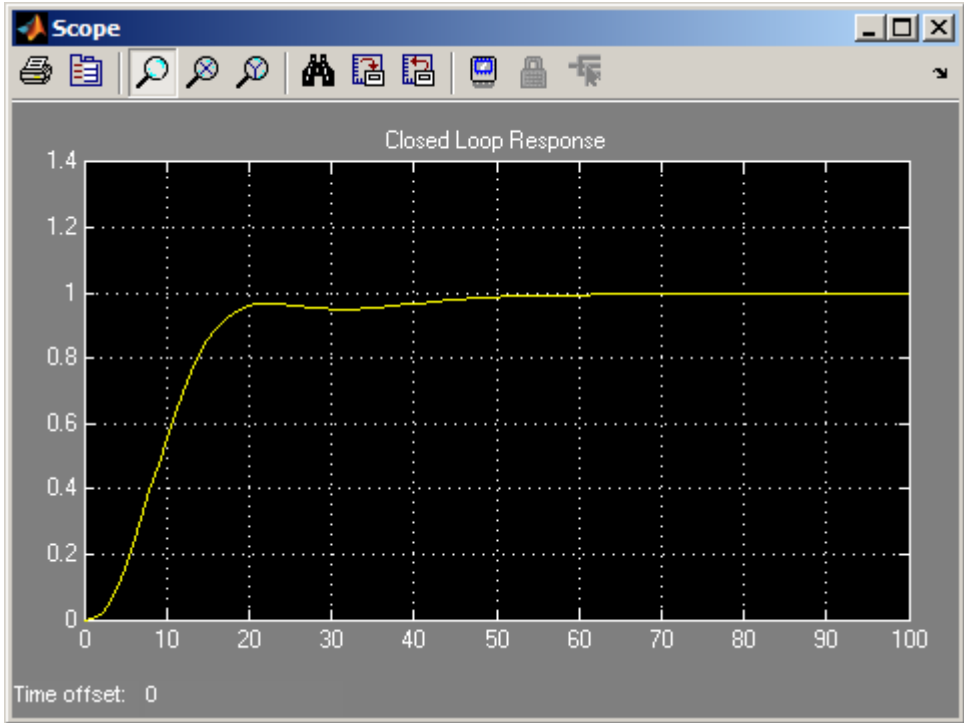
```
Kp =
    0.5894
```

```
Ki =
    0.0605
```

```
Kd =
    5.5295
```

The last value in the **Objective value** column of the output shows that the maximum value for all the time steps is 0.9997. The closed loop response with this result is shown in the figure *Closed-Loop Response Using fminimax* on page 6-245.

This solution differs from the solution obtained in “lsqnonlin with a Simulink Model” on page 6-209 because you are solving different problem formulations.



Closed-Loop Response Using fminimax

Signal Processing Using `fgoalattain`

Consider designing a linear-phase Finite Impulse Response (FIR) filter. The problem is to design a lowpass filter with magnitude one at all frequencies between 0 and 0.1 Hz and magnitude zero between 0.15 and 0.5 Hz.

The frequency response $H(f)$ for such a filter is defined by

$$\begin{aligned} H(f) &= \sum_{n=0}^{2M} h(n)e^{-j2\pi fn} \\ &= A(f)e^{-j2\pi fM}, \\ A(f) &= \sum_{n=0}^{M-1} a(n)\cos(2\pi fn), \end{aligned} \tag{6-119}$$

where $A(f)$ is the magnitude of the frequency response. One solution is to apply a goal attainment method to the magnitude of the frequency response. Given a function that computes the magnitude, `fgoalattain` will attempt to vary the magnitude coefficients $a(n)$ until the magnitude response matches the desired response within some tolerance. The function that computes the magnitude response is given in `filtmin.m`. This function uses `a`, the magnitude function coefficients, and `w`, the discretization of the frequency domain of interest.

To set up a goal attainment problem, you must specify the `goal` and `weights` for the problem. For frequencies between 0 and 0.1, the goal is one. For frequencies between 0.15 and 0.5, the goal is zero. Frequencies between 0.1 and 0.15 are not specified, so no goals or weights are needed in this range.

This information is stored in the variable `goal` passed to `fgoalattain`. The length of `goal` is the same as the length returned by the function `filtmin`. So that the goals are equally satisfied, usually `weight` would be set to `abs(goal)`. However, since some of the goals are zero, the effect of using `weight=abs(goal)` will force the objectives with `weight 0` to be satisfied as hard constraints, and the objectives with `weight 1` possibly to be underattained (see “Goal Attainment Method” on page 6-237). Because all the goals are close in magnitude, using a `weight` of unity for all goals will give them equal priority. (Using `abs(goal)` for the weights is more important when the magnitude of `goal` differs more significantly.) Also, setting


```
options = optimoptions('fgoalattain','GoalsExactAchieve',length(goal));
```

specifies that each objective should be as near as possible to its goal value (neither greater nor less than).

Step 1: Write a file filtmin.m

```
function y = filtmin(a,w)
n = length(a);
y = cos(w*(0:n-1)*2*pi)*a ;
```

Step 2: Invoke optimization routine

```
% Plot with initial coefficients
a0 = ones(15,1);
incr = 50;
w = linspace(0,0.5,incr);

y0 = filtmin(a0,w);
clf, plot(w,y0,'-.b');
drawnow;


% Set up the goal attainment problem
w1 = linspace(0,0.1,incr) ;
w2 = linspace(0.15,0.5,incr);
w0 = [w1 w2];
goal = [1.0*ones(1,length(w1)) zeros(1,length(w2))];
weight = ones(size(goal));

% Call fgoalattain
options = optimoptions('fgoalattain','GoalsExactAchieve',length(goal));
[a,fval,attainfactor,exitflag]=fgoalattain(@(x)filtmin(x,w0),...
    a0,goal,weight,[],[],[],[],[],[],[],[],options);

% Plot with the optimized (final) coefficients
y = filtmin(a,w);
hold on, plot(w,y,'r')
axis([0 0.5 -3 3])
xlabel('Frequency (Hz)')
ylabel('Magnitude Response (dB)')
```

```
legend('initial', 'final')  
grid on
```

Compare the magnitude response computed with the initial coefficients and the final coefficients (Magnitude Response with Initial and Final Magnitude Coefficients on page 6-248). Note that you could use the `firpm` function in Signal Processing Toolbox™ software to design this filter.



Magnitude Response with Initial and Final Magnitude Coefficients

Equation Solving Algorithms

In this section...

“Equation Solving Definition” on page 6-249

“Trust-Region Reflective `fsolve` Algorithm” on page 6-249

“Trust-Region Dogleg Method” on page 6-252

“Levenberg-Marquardt Method” on page 6-255

“`\` Algorithm” on page 6-255

“`fzero` Algorithm” on page 6-255

Equation Solving Definition

Given a set of n nonlinear functions $F_i(x)$, where n is the number of components of the vector x , the goal of equation solving is to find a vector x that makes all $F_i(x) = 0$.

`fsolve` attempts to solve systems of equations by minimizing the sum of squares of the components. If the sum of squares is zero, the system of equation is solved. `fsolve` has three algorithms:

- Trust-region-reflective
- Trust-region dogleg
- Levenberg-Marquardt

All algorithms are large-scale; see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12.

The `fzero` function solves a single one-dimensional equation.

The `\` function solves systems of linear equations.

Trust-Region Reflective `fsolve` Algorithm

Many of the methods used in Optimization Toolbox solvers are based on *trust regions*, a simple yet powerful concept in optimization.

To understand the trust-region approach to optimization, consider the unconstrained minimization problem, minimize $f(x)$, where the function takes vector arguments and returns scalars. Suppose you are at a point x in n -space and you want to improve, i.e., move to a point with a lower function value. The basic idea is to approximate f with a simpler function q , which reasonably reflects the behavior of function f in a neighborhood N around the point x . This neighborhood is the trust region. A trial step s is computed by minimizing (or approximately minimizing) over N . This is the trust-region subproblem,

$$\min_s \{q(s), s \in N\}. \quad (6-120)$$

The current point is updated to be $x + s$ if $f(x + s) < f(x)$; otherwise, the current point remains unchanged and N , the region of trust, is shrunk and the trial step computation is repeated.

The key questions in defining a specific trust-region approach to minimizing $f(x)$ are how to choose and compute the approximation q (defined at the current point x), how to choose and modify the trust region N , and how accurately to solve the trust-region subproblem. This section focuses on the unconstrained problem. Later sections discuss additional complications due to the presence of constraints on the variables.

In the standard trust-region method ([48]), the quadratic approximation q is defined by the first two terms of the Taylor approximation to F at x ; the neighborhood N is usually spherical or ellipsoidal in shape. Mathematically the trust-region subproblem is typically stated

$$\min \left\{ \frac{1}{2} s^T H s + s^T g \text{ such that } \|Ds\| \leq \Delta \right\}, \quad (6-121)$$

where g is the gradient of f at the current point x , H is the Hessian matrix (the symmetric matrix of second derivatives), D is a diagonal scaling matrix, Δ is a positive scalar, and $\| \cdot \|$ is the 2-norm. Good algorithms exist for solving Equation 6-121 (see [48]); such algorithms typically involve the computation of a full eigensystem and a Newton process applied to the secular equation

$$\frac{1}{\Delta} - \frac{1}{\|s\|} = 0.$$

Such algorithms provide an accurate solution to Equation 6-121. However, they require time proportional to several factorizations of H . Therefore, for trust-region problems a different approach is needed. Several approximation and heuristic strategies, based on Equation 6-121, have been proposed in the literature ([42] and [50]). The approximation approach followed in Optimization Toolbox solvers is to restrict the trust-region subproblem to a two-dimensional subspace S ([39] and [42]). Once the subspace S has been computed, the work to solve Equation 6-121 is trivial even if full eigenvalue/eigenvector information is needed (since in the subspace, the problem is only two-dimensional). The dominant work has now shifted to the determination of the subspace.

The two-dimensional subspace S is determined with the aid of a preconditioned conjugate gradient process described below. The solver defines S as the linear space spanned by s_1 and s_2 , where s_1 is in the direction of the gradient g , and s_2 is either an approximate Newton direction, i.e., a solution to

$$H \cdot s_2 = -g, \quad (6-122)$$

or a direction of negative curvature,

$$s_2^T \cdot H \cdot s_2 < 0. \quad (6-123)$$

The philosophy behind this choice of S is to force global convergence (via the steepest descent direction or negative curvature direction) and achieve fast local convergence (via the Newton step, when it exists).

A sketch of unconstrained minimization using trust-region ideas is now easy to give:

- 1** Formulate the two-dimensional trust-region subproblem.
- 2** Solve Equation 6-121 to determine the trial step s .
- 3** If $f(x + s) < f(x)$, then $x = x + s$.
- 4** Adjust Δ .

These four steps are repeated until convergence. The trust-region dimension Δ is adjusted according to standard rules. In particular, it is decreased if the

trial step is not accepted, i.e., $f(x + s) \geq f(x)$. See [46] and [49] for a discussion of this aspect.

Optimization Toolbox solvers treat a few important special cases of f with specialized functions: nonlinear least-squares, quadratic functions, and linear least-squares. However, the underlying algorithmic ideas are the same as for the general case. These special cases are discussed in later sections.

Preconditioned Conjugate Gradient Method

A popular way to solve large symmetric positive definite systems of linear equations $Hp = -g$ is the method of Preconditioned Conjugate Gradients (PCG). This iterative approach requires the ability to calculate matrix-vector products of the form Hv where v is an arbitrary vector. The symmetric positive definite matrix M is a *preconditioner* for H . That is, $M = C^2$, where $C^{-1}HC^{-1}$ is a well-conditioned matrix or a matrix with clustered eigenvalues.

In a minimization context, you can assume that the Hessian matrix H is symmetric. However, H is guaranteed to be positive definite only in the neighborhood of a strong minimizer. Algorithm PCG exits when a direction of negative (or zero) curvature is encountered, i.e., $d^T Hd \leq 0$. The PCG output direction, p , is either a direction of negative curvature or an approximate (*tol* controls how approximate) solution to the Newton system $Hp = -g$. In either case p is used to help define the two-dimensional subspace used in the trust-region approach discussed in “Trust-Region Methods for Nonlinear Minimization” on page 6-5.

Trust-Region Dogleg Method

Another approach is to solve a linear system of equations to find the search direction, namely, Newton’s method says to solve for the search direction d_k such that

$$\begin{aligned} J(x_k)d_k &= -F(x_k) \\ x_{k+1} &= x_k + d_k, \end{aligned}$$

where $J(x_k)$ is the n -by- n Jacobian

$$J(x_k) = \begin{bmatrix} \nabla F_1(x_k)^T \\ \nabla F_2(x_k)^T \\ \vdots \\ \nabla F_n(x_k)^T \end{bmatrix}.$$

Newton's method can run into difficulties. $J(x_k)$ may be singular, and so the Newton step d_k is not even defined. Also, the exact Newton step d_k may be expensive to compute. In addition, Newton's method may not converge if the starting point is far from the solution.

Using trust-region techniques (introduced in "Trust-Region Methods for Nonlinear Minimization" on page 6-5) improves robustness when starting far from the solution and handles the case when $J(x_k)$ is singular. To use a trust-region strategy, a merit function is needed to decide if x_{k+1} is better or worse than x_k . A possible choice is

$$\min_d f(d) = \frac{1}{2} F(x_k + d)^T F(x_k + d).$$

But a minimum of $f(d)$ is not necessarily a root of $F(x)$.

The Newton step d_k is a root of

$$M(x_k + d) = F(x_k) + J(x_k)d,$$

and so it is also a minimum of $m(d)$, where

$$\begin{aligned} \min_d m(d) &= \frac{1}{2} \|M(x_k + d)\|_2^2 = \frac{1}{2} \|F(x_k) + J(x_k)d\|_2^2 \\ &= \frac{1}{2} F(x_k)^T F(x_k) + d^T J(x_k)^T F(x_k) + \frac{1}{2} d^T J(x_k)^T J(x_k)d \end{aligned} \quad (6-124)$$

Then $m(d)$ is a better choice of merit function than $f(d)$, and so the trust-region subproblem is

$$\min_d \left[\frac{1}{2} F(x_k)^T F(x_k) + d^T J(x_k)^T F(x_k) + \frac{1}{2} d^T J(x_k)^T J(x_k) d \right], \quad (6-125)$$

such that $\|D \cdot d\| \leq \Delta$. This subproblem can be efficiently solved using a dogleg strategy.

For an overview of trust-region methods, see Conn [4], and Nocedal [31].

Trust-Region Dogleg Implementation

The key feature of this algorithm is the use of the Powell dogleg procedure for computing the step d , which minimizes Equation 6-125. For a detailed description, see Powell [34].

The step d is constructed from a convex combination of a Cauchy step (a step along the steepest descent direction) and a Gauss-Newton step for $f(x)$. The Cauchy step is calculated as

$$d_C = -\alpha J(x_k)^T F(x_k),$$

where α is chosen to minimize Equation 6-124.

The Gauss-Newton step is calculated by solving

$$J(x_k) \cdot d_{GN} = -F(x_k),$$

using the MATLAB \ (matrix left division) operator.

The step d is chosen so that

$$d = d_C + \lambda(d_{GN} - d_C),$$

where λ is the largest value in the interval $[0,1]$ such that $\|d\| \leq \Delta$. If J_k is (nearly) singular, d is just the Cauchy direction.

The dogleg algorithm is efficient since it requires only one linear solve per iteration (for the computation of the Gauss-Newton step). Additionally, it can be more robust than using the Gauss-Newton method with a line search.

Levenberg-Marquardt Method

The Levenberg-Marquardt [25], and [27] method uses a search direction that is a solution of the linear set of equations

$$\left(J(x_k)^T J(x_k) + \lambda_k I \right) d_k = -J(x_k)^T F(x_k), \quad (6-126)$$

or, optionally, of the equations

$$\left(J(x_k)^T J(x_k) + \lambda_k \text{diag} \left(J(x_k)^T J(x_k) \right) \right) d_k = -J(x_k)^T F(x_k), \quad (6-127)$$

where the scalar λ_k controls both the magnitude and direction of d_k . Set option `ScaleProblem` to 'none' to choose Equation 6-126, and set `ScaleProblem` to 'Jacobian' to choose Equation 6-127.

When λ_k is zero, the direction d_k is the Gauss-Newton method. As λ_k tends to infinity, d_k tends towards the steepest descent direction, with magnitude tending to zero. This implies that for some sufficiently large λ_k , the term $F(x_k + d_k) < F(x_k)$ holds true. The term λ_k can therefore be controlled to ensure descent even when second-order terms, which restrict the efficiency of the Gauss-Newton method, are encountered. The Levenberg-Marquardt method therefore uses a search direction that is a cross between the Gauss-Newton direction and the steepest descent direction.

\ Algorithm

This algorithm is described in the MATLAB arithmetic operators section for `\ (mldivide)`.

fzero Algorithm

`fzero` attempts to find the root of a scalar function f of a scalar variable x .

`fzero` looks for an interval around an initial point such that $f(x)$ changes sign. If you give an initial interval instead of an initial point, `fzero` checks to make sure $f(x)$ has different signs at the endpoints of the interval. The initial interval must be finite; it cannot contain $\pm\text{Inf}$.

`fzero` uses a combination of interval bisection, linear interpolation, and inverse quadratic interpolation in order to locate a root of $f(x)$. See `fzero` for more information.

Nonlinear Equations with Analytic Jacobian

This example demonstrates the use of the default trust-region-dogleg `fsolve` algorithm (see “Large-Scale vs. Medium-Scale Algorithms” on page 2-12). It is intended for problems where

- The system of nonlinear equations is square, i.e., the number of equations equals the number of unknowns.
- There exists a solution x such that $F(x) = 0$.

The example uses `fsolve` to obtain the minimum of the banana (or Rosenbrock) function by deriving and then solving an equivalent system of nonlinear equations. The Rosenbrock function, which has a minimum of $F(x) = 0$, is a common test problem in optimization. It has a high degree of nonlinearity and converges extremely slowly if you try to use steepest descent type methods. It is given by

$$f(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2.$$

First generalize this function to an n -dimensional function, for any positive, even value of n :

$$f(x) = \sum_{i=1}^{n/2} 100(x_{2i} - x_{2i-1}^2)^2 + (1 - x_{2i-1})^2.$$

This function is referred to as the generalized Rosenbrock function. It consists of n squared terms involving n unknowns.

Before you can use `fsolve` to find the values of x such that $F(x) = 0$, i.e., obtain the minimum of the generalized Rosenbrock function, you must rewrite the function as the following equivalent system of nonlinear equations:

$$\begin{aligned}
 F(1) &= 1 - x_1 \\
 F(2) &= 10(x_2 - x_1^2) \\
 F(3) &= 1 - x_3 \\
 F(4) &= 10(x_4 - x_3^2) \\
 &\vdots \\
 F(n-1) &= 1 - x_{n-1} \\
 F(n) &= 10(x_n - x_{n-1}^2).
 \end{aligned}$$

This system is square, and you can use `fsolve` to solve it. As the example demonstrates, this system has a unique solution given by $x_i = 1$, $i = 1, \dots, n$.

Step 1: Write a file `bananaobj.m` to compute the objective function values and the Jacobian.

```

function [F,J] = bananaobj(x)
% Evaluate the vector function and the Jacobian matrix for
% the system of nonlinear equations derived from the general
% n-dimensional Rosenbrock function.
% Get the problem size
n = length(x);
if n == 0, error('Input vector, x, is empty.');
```

```

end
if mod(n,2) ~= 0,
    error('Input vector, x ,must have an even number of components.');
```

```

end
% Evaluate the vector function
odds = 1:2:n;
evens = 2:2:n;
F = zeros(n,1);
F(odds,1) = 1-x(odds);
F(evens,1) = 10.*(x(evens)-x(odds).^2);
% Evaluate the Jacobian matrix if nargin > 1
if nargin > 1
    c = -ones(n/2,1);    C = sparse(odds,odds,c,n,n);
    d = 10*ones(n/2,1); D = sparse(evens,evens,d,n,n);
    e = -20.*x(odds);   E = sparse(evens,odds,e,n,n);

```

```
J = C + D + E;
end
```

Step 2: Call the solve routine for the system of equations.

```
n = 64;
x0(1:n,1) = -1.9;
x0(2:2:n,1) = 2;
options = optimoptions(@fsolve,'Display','iter','Jacobian','on');
[x,F,exitflag,output,JAC] = fsolve(@bananaobj,x0,options);
```

Use the starting point $x(i) = -1.9$ for the odd indices, and $x(i) = 2$ for the even indices. Set `Display` to `'iter'` to see the solver's progress. Set `Jacobian` to `'on'` to use the Jacobian defined in `bananaobj.m`. The `fsolve` function generates the following output:

Iteration	Func-count	f(x)	Norm of step	First-order optimality	Trust-region radius
0	1	8563.84		615	1
1	2	3093.71	1	329	1
2	3	225.104	2.5	34.8	2.5
3	4	212.48	6.25	34.1	6.25
4	5	212.48	6.25	34.1	6.25
5	6	102.771	1.5625	6.39	1.56
6	7	102.771	3.90625	6.39	3.91
7	8	87.7443	0.976563	2.19	0.977
8	9	74.1426	2.44141	6.27	2.44
9	10	74.1426	2.44141	6.27	2.44
10	11	52.497	0.610352	1.52	0.61
11	12	41.3297	1.52588	4.63	1.53
12	13	34.5115	1.52588	6.97	1.53
13	14	16.9716	1.52588	4.69	1.53
14	15	8.16797	1.52588	3.77	1.53
15	16	3.55178	1.52588	3.56	1.53
16	17	1.38476	1.52588	3.31	1.53
17	18	0.219553	1.16206	1.66	1.53
18	19	0	0.0468565	0	1.53

Equation solved.

fsolve completed because the vector of function values is near zero as measured by the default value of the function tolerance, and the problem appears regular as measured by the gradient.

Nonlinear Equations with Finite-Difference Jacobian

In the example “Nonlinear Equations with Analytic Jacobian” on page 6-257, the function `bananaobj` evaluates F and computes the Jacobian J . What if the code to compute the Jacobian is not available? By default, if you do not indicate that the Jacobian can be computed in the objective function (by setting the `Jacobian` option in `options` to `'on'`), `fsolve`, `lsqnonlin`, and `lsqcurvefit` instead use finite differencing to approximate the Jacobian. This is the default Jacobian option. You can select finite differencing by setting `Jacobian` to `'off'` using `optimoptions`.

This example uses `bananaobj` from the example “Nonlinear Equations with Analytic Jacobian” on page 6-257 as the objective function, but sets `Jacobian` to `'off'` so that `fsolve` approximates the Jacobian and ignores the second `bananaobj` output.

```
n = 64;
x0(1:n,1) = -1.9;
x0(2:2:n,1) = 2;
options = optimoptions(@fsolve,'Display','iter','Jacobian','off');
[x,F,exitflag,output,JAC] = fsolve(@bananaobj,x0,options);
```

The example produces the following output:

Iteration	Func-count	f(x)	Norm of step	First-order optimality	Trust-region radius
0	65	8563.84		615	1
1	130	3093.71	1	329	1
2	195	225.104	2.5	34.8	2.5
3	260	212.48	6.25	34.1	6.25
4	261	212.48	6.25	34.1	6.25
5	326	102.771	1.5625	6.39	1.56
6	327	102.771	3.90625	6.39	3.91
7	392	87.7443	0.976562	2.19	0.977
8	457	74.1426	2.44141	6.27	2.44
9	458	74.1426	2.44141	6.27	2.44
10	523	52.497	0.610352	1.52	0.61
11	588	41.3297	1.52588	4.63	1.53
12	653	34.5115	1.52588	6.97	1.53
13	718	16.9716	1.52588	4.69	1.53

14	783	8.16797	1.52588	3.77	1.53
15	848	3.55178	1.52588	3.56	1.53
16	913	1.38476	1.52588	3.31	1.53
17	978	0.219553	1.16206	1.66	1.53
18	1043	0	0.0468565	0	1.53

Equation solved.

fsolve completed because the vector of function values is near zero as measured by the default value of the function tolerance, and the problem appears regular as measured by the gradient.

The finite-difference version of this example requires the same number of iterations to converge as the analytic Jacobian version in the preceding example. It is generally the case that both versions converge at about the same rate in terms of iterations. However, the finite-difference version requires many additional function evaluations. The cost of these extra evaluations might or might not be significant, depending on the particular problem.

Nonlinear Equations with Jacobian

Consider the problem of finding a solution to a system of nonlinear equations whose Jacobian is sparse. The dimension of the problem in this example is 1000. The goal is to find x such that $F(x) = 0$. Assuming $n = 1000$, the nonlinear equations are

$$F(1) = 3x_1 - 2x_1^2 - 2x_2 + 1,$$

$$F(i) = 3x_i - 2x_i^2 - x_{i-1} - 2x_{i+1} + 1,$$

$$F(n) = 3x_n - 2x_n^2 - x_{n-1} + 1.$$

To solve a large nonlinear system of equations, $F(x) = 0$, you can use the trust-region reflective algorithm available in `fsolve`, a large-scale algorithm (“Large-Scale vs. Medium-Scale Algorithms” on page 2-12).

Step 1: Write a file `nlsf1.m` that computes the objective function values and the Jacobian.

```
function [F,J] = nlsf1(x)
% Evaluate the vector function
n = length(x);
F = zeros(n,1);
i = 2:(n-1);
F(i) = (3-2*x(i)).*x(i)-x(i-1)-2*x(i+1) + 1;
F(n) = (3-2*x(n)).*x(n)-x(n-1) + 1;
F(1) = (3-2*x(1)).*x(1)-2*x(2) + 1;
% Evaluate the Jacobian if nargout > 1
if nargout > 1
    d = -4*x + 3*ones(n,1); D = sparse(1:n,1:n,d,n,n);
    c = -2*ones(n-1,1); C = sparse(1:n-1,2:n,c,n,n);
    e = -ones(n-1,1); E = sparse(2:n,1:n-1,e,n,n);
    J = C + D + E;
end
```

Step 2: Call the solve routine for the system of equations.

```
xstart = -ones(1000,1);
fun = @nlsf1;
options = optimoptions(@fsolve,'Display','iter',...
    'Algorithm','trust-region-reflective',...
    'Jacobian','on','PrecondBandWidth',0);
[x,fval,exitflag,output] = fsolve(fun,xstart,options);
```

A starting point is given as well as the function name. The default method for `fsolve` is `trust-region-dogleg`, so it is necessary to specify `'Algorithm'` as `'trust-region-reflective'` in the `options` argument in order to select the `trust-region-reflective` algorithm. Setting the `Display` option to `'iter'` causes `fsolve` to display the output at each iteration. Setting `Jacobian` to `'on'`, causes `fsolve` to use the Jacobian information available in `nlsf1.m`.

The commands display this output:

Iteration	Func-count	f(x)	Norm of step	First-order optimality
0	1	1011		19
1	2	16.1942	7.91898	2.35
2	3	0.0228027	1.33142	0.291
3	4	0.000103359	0.0433329	0.0201
4	5	7.3792e-07	0.0022606	0.000946
5	6	4.02299e-10	0.000268381	4.12e-05

Equation solved, inaccuracy possible.

The vector of function values is near zero, as measured by the default value of the function tolerance. However, the last step was ineffective.

A linear system is (approximately) solved in each major iteration using the preconditioned conjugate gradient method. Setting `PrecondBandWidth` to 0 in `options` means a diagonal preconditioner is used. (`PrecondBandWidth` specifies the bandwidth of the preconditioning matrix. A bandwidth of 0 means there is only one diagonal in the matrix.)

From the first-order optimality values, fast linear convergence occurs. The number of conjugate gradient (CG) iterations required per major iteration is low, at most five for a problem of 1000 dimensions, implying that the linear systems are not very difficult to solve in this case (though more work is required as convergence progresses).

If you want to use a tridiagonal preconditioner, i.e., a preconditioning matrix with three diagonals (or bandwidth of one), set `PrecondBandWidth` to the value 1:

```
options = optimoptions(@fsolve,'Display','iter','Jacobian','on',...
    'Algorithm','trust-region-reflective','PrecondBandWidth',1);
[x,fval,exitflag,output] = fsolve(fun,xstart,options);
```

In this case the output is

Iteration	Func-count	f(x)	Norm of step	First-order optimality
0	1	1011		19
1	2	16.0839	7.92496	1.92
2	3	0.0458181	1.3279	0.579
3	4	0.000101184	0.0631898	0.0203
4	5	3.16615e-07	0.00273698	0.00079
5	6	9.72481e-10	0.00018111	5.82e-05

Equation solved, inaccuracy possible.

The vector of function values is near zero, as measured by the default value of the function tolerance. However, the last step was ineffective.

Note that although the same number of iterations takes place, the number of PCG iterations has dropped, so less work is being done per iteration. See “Preconditioned Conjugate Gradient Method” on page 6-29.

Setting `PrecondBandWidth` to `Inf` (this is the default) means that the solver uses Cholesky factorization rather than PCG.

Nonlinear Equations with Jacobian Sparsity Pattern

In the example “Nonlinear Equations with Analytic Jacobian” on page 6-257, the function `nlsf1` computes the Jacobian J , a sparse matrix, along with the evaluation of F . What if the code to compute the Jacobian is not available? By default, if you do not indicate that the Jacobian can be computed in `nlsf1` (by setting the `Jacobian` option in `options` to 'on'), `fsolve`, `lsqnonlin`, and `lsqcurvefit` instead uses finite differencing to approximate the Jacobian.

In order for this finite differencing to be as efficient as possible, you should supply the sparsity pattern of the Jacobian, by setting `JacobPattern` to a sparse matrix `Jstr` in `options`. That is, supply a sparse matrix `Jstr` whose nonzero entries correspond to nonzeros of the Jacobian for all x . Indeed, the nonzeros of `Jstr` can correspond to a superset of the nonzero locations of J ; however, in general the computational cost of the sparse finite-difference procedure will increase with the number of nonzeros of `Jstr`.

Providing the sparsity pattern can drastically reduce the time needed to compute the finite differencing on large problems. If the sparsity pattern is not provided (and the Jacobian is not computed in the objective function either) then, in this problem with 1000 variables, the finite-differencing code attempts to compute all 1000-by-1000 entries in the Jacobian. But in this case there are only 2998 nonzeros, substantially less than the 1,000,000 possible nonzeros the finite-differencing code attempts to compute. In other words, this problem is solvable if you provide the sparsity pattern. If not, most computers run out of memory when the full dense finite-differencing is attempted. On most small problems, it is not essential to provide the sparsity structure.

Suppose the sparse matrix `Jstr`, computed previously, has been saved in file `nlsdat1.mat`. The following driver calls `fsolve` applied to `nlsf1a`, which is `nlsf1` without the Jacobian. Sparse finite-differencing is used to estimate the sparse Jacobian matrix as needed.

Step 1: Write a file `nlsf1a.m` that computes the objective function values.

```
function F = nlsf1a(x)
% Evaluate the vector function
n = length(x);
```

```
F = zeros(n,1);
i = 2:(n-1);
F(i) = (3-2*x(i)).*x(i)-x(i-1)-2*x(i+1) + 1;
F(n) = (3-2*x(n)).*x(n)-x(n-1) + 1;
F(1) = (3-2*x(1)).*x(1)-2*x(2) + 1;
```

Step 2: Call the system of equations solve routine.

```
xstart = -ones(1000,1);
fun = @nlsf1a;
load nlsdat1 % Get Jstr
options = optimoptions(@fsolve,'Display','iter','JacobPattern',Jstr,...
    'Algorithm','trust-region-reflective','PrecondBandWidth',1);
[x,fval,exitflag,output] = fsolve(fun,xstart,options);
```

In this case, the output displayed is

Iteration	Func-count	f(x)	Norm of step	First-order optimality
0	5	1011		19
1	10	16.0839	7.92496	1.92
2	15	0.0458179	1.3279	0.579
3	20	0.000101184	0.0631896	0.0203
4	25	3.16616e-07	0.00273698	0.00079
5	30	9.72483e-10	0.00018111	5.82e-05

Equation solved, inaccuracy possible.

The vector of function values is near zero, as measured by the default value of the function tolerance. However, the last step was ineffective.

Alternatively, it is possible to choose a sparse direct linear solver (i.e., a sparse QR factorization) by indicating a “complete” preconditioner. For example, if you set `PrecondBandWidth` to `Inf`, then a sparse direct linear solver is used instead of a preconditioned conjugate gradient iteration:

```
xstart = -ones(1000,1);
fun = @nlsf1a;
load nlsdat1 % Get Jstr
options = optimoptions(@fsolve,'Display','iter','JacobPattern',Jstr,...
    'Algorithm','trust-region-reflective','PrecondBandWidth',inf);
```

```
[x,fval,exitflag,output] = fsolve(fun,xstart,options);
```

and the resulting display is

Iteration	Func-count	f(x)	Norm of step	First-order optimality
0	5	1011		19
1	10	15.9018	7.92421	1.89
2	15	0.0128161	1.32542	0.0746
3	20	1.73502e-08	0.0397923	0.000196
4	25	1.10716e-18	4.55495e-05	2.74e-09

Equation solved.

fsolve completed because the vector of function values is near zero as measured by the default value of the function tolerance, and the problem appears regular as measured by the gradient.

When the sparse direct solvers are used, the CG iteration is 0 for that (major) iteration, as shown in the output under **CG-Iterations**. Notice that the final optimality and $f(x)$ value (which for `fsolve`, $f(x)$, is the sum of the squares of the function values) are closer to zero than using the PCG method, which is often the case.

Nonlinear Systems with Constraints

In this section...

“Solve Equations with Inequality Constraints” on page 6-269

“Use Different Start Points” on page 6-270

“Use Different Algorithms” on page 6-270

“Use lsqnonlin with Bounds” on page 6-271

“Set Equations and Inequalities as fmincon Constraints” on page 6-272

Solve Equations with Inequality Constraints

`fsolve` solves systems of nonlinear equations. However, it does not allow you to include any constraints, even bound constraints. The question is, how can you solve systems of nonlinear equations when you have constraints?

The short answer is, there are no guarantees that a solution exists that satisfies your constraints. There is no guarantee that any solution exists, even one that does not satisfy your constraints. Nevertheless, there are techniques that can help you search for solutions that satisfy your constraints.

To illustrate the techniques, consider how to solve the equations

$$\begin{aligned}
 F_1(x) &= (x_1 + 1)(10 - x_1) \frac{1 + x_2^2}{1 + x_2^2 + x_2} \\
 F_2(x) &= (x_2 + 2)(20 - x_2) \frac{1 + x_1^2}{1 + x_1^2 + x_1},
 \end{aligned} \tag{6-128}$$

where the components of x must be nonnegative. Clearly, there are four solutions to the equations:

$$\begin{aligned}
 x &= (-1, -2) \\
 x &= (10, -2), \\
 x &= (-1, 20), \\
 x &= (10, 20).
 \end{aligned}$$

There is only one solution that satisfies the constraints, namely $x = (10,20)$.

To solve the equations numerically, first enter code to calculate $F(x)$.

```
function F = fbnd(x)

F(1) = (x(1)+1)*(10-x(1))*(1+x(2)^2)/(1+x(2)^2+x(2));
F(2) = (x(2)+2)*(20-x(2))*(1+x(1)^2)/(1+x(1)^2+x(1));
```

Save this code as the file `fbnd.m` on your MATLAB path.

Use Different Start Points

Generally, a system of N equations in N variables has isolated solutions, meaning each solution has no nearby neighbors that are also solutions. So one way to search for a solution that satisfies some constraints is to generate a number of initial points x_0 , and run `fsolve` starting at each x_0 .

For this example, to look for a solution to Equation 6-128, take 10 random points that are normally distributed with mean 0 and standard deviation 100.

```
N = 10; % try 10 random start points
pts = 100*randn(N,2); % initial points are rows in pts
soln = zeros(N,2); % allocate solution
opts = optimoptions('fsolve','Display','off');
rng('default') % for reproducibility
for k = 1:N
    soln(k,:) = fsolve(@fbnd,pts(k,:),opts); % find solutions
end
```

Examine the solutions in `soln`, and you find several that satisfy the constraints.

Use Different Algorithms

There are three `fsolve` algorithms. Each can lead to different solutions.

For this example, take $x_0 = [1,9]$ and examine the solution each algorithm returns.


```
x0 = [1,9];
opts = optimoptions(@fsolve,'Display','off',...
    'Algorithm','trust-region-dogleg');
x1 = fsolve(@fbnd,x0,opts)
```

```
x1 =
    -1.0000    -2.0000
```

```
opts.Algorithm = 'trust-region-reflective';
x2 = fsolve(@fbnd,x0,opts)
```

```
x2 =
    -1.0000    20.0000
```

```
opts.Algorithm = 'levenberg-marquardt';
x3 = fsolve(@fbnd,x0,opts)
```

```
x3 =
    0.9523    8.9941
```

Here, all three algorithms find different solutions for the same initial point. In fact, x_3 is not even a solution, but is simply a locally stationary point.

Use lsqnonlin with Bounds

`lsqnonlin` tries to minimize the sum of squares of the components of a vector function $F(x)$. Therefore, it attempts to solve the equation $F(x) = 0$. Furthermore, `lsqnonlin` accepts bound constraints.

Formulate the example problem for `lsqnonlin` and solve it.

```
lb = [0,0];
rng('default')
```

```
x0 = 100*randn(2,1);  
[x res] = lsqnonlin(@fbnd,x0,lb)
```

```
x =
```

```
10.0000  
20.0000
```

```
res =
```

```
2.4783e-25
```

You can use `lsqnonlin` with the Global Optimization Toolbox `MultiStart` solver to search over many initial points automatically. See “MultiStart Using `lsqcurvefit` or `lsqnonlin`”.

Set Equations and Inequalities as `fmincon` Constraints

You can reformulate the problem and use `fmincon` as follows:

- Give a constant objective function, such as $@(x)0$, which evaluates to 0 for each x .
- Set the `fsolve` objective function as the nonlinear equality constraints in `fmincon`.
- Give any other constraints in the usual `fmincon` syntax.

For this example, write a function file for the nonlinear inequality constraint.

```
function [c,ceq] = fminconstr(x)
```

```
c = []; % no nonlinear inequality
```

```
ceq = fbnd(x); % the fsolve objective is fmincon constraints
```

Save this code as the file `fminconstr.m` on your MATLAB path.

Solve the constrained problem.

```
lb = [0,0]; % lower bound constraint
rng('default') % reproducible initial point
x0 = 100*randn(2,1);
opts = optimoptions(@fmincon,'Algorithm','interior-point','Display','off');
x = fmincon(@(x)0,x0,[],[],[],[],lb,[],@fminconstr,opts)
```

```
x =
```

```
10.0000
20.0000
```

Bibliography

- [1] Biggs, M.C., “Constrained Minimization Using Recursive Quadratic Programming,” *Towards Global Optimization* (L.C.W. Dixon and G.P. Szergo, eds.), North-Holland, pp 341–349, 1975.
- [2] Brayton, R.K., S.W. Director, G.D. Hachtel, and L. Vidigal, “A New Algorithm for Statistical Circuit Design Based on Quasi-Newton Methods and Function Splitting,” *IEEE Transactions on Circuits and Systems*, Vol. CAS-26, pp 784–794, Sept. 1979.
- [3] Broyden, C.G., “The Convergence of a Class of Double-rank Minimization Algorithms,”; *J. Inst. Maths. Applics.*, Vol. 6, pp 76–90, 1970.
- [4] Conn, N.R., N.I.M. Gould, and Ph.L. Toint, *Trust-Region Methods*, MPS/SIAM Series on Optimization, SIAM and MPS, 2000.
- [5] Dantzig, G., *Linear Programming and Extensions*, Princeton University Press, Princeton, 1963.
- [6] Dantzig, G.B., A. Orden, and P. Wolfe, “Generalized Simplex Method for Minimizing a Linear Form Under Linear Inequality Restraints,” *Pacific Journal Math.*, Vol. 5, pp. 183–195, 1955.
- [7] Davidon, W.C., “Variable Metric Method for Minimization,” *A.E.C. Research and Development Report*, ANL-5990, 1959.
- [8] Dennis, J.E., Jr., “Nonlinear least-squares,” *State of the Art in Numerical Analysis* ed. D. Jacobs, Academic Press, pp 269–312, 1977.
- [9] Dennis, J.E., Jr. and R.B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall Series in Computational Mathematics, Prentice-Hall, 1983.
- [10] Fleming, P.J., “Application of Multiobjective Optimization to Compensator Design for SISO Control Systems,” *Electronics Letters*, Vol. 22, No. 5, pp 258–259, 1986.

- [11] Fleming, P.J., "Computer-Aided Control System Design of Regulators using a Multiobjective Optimization Approach," *Proc. IFAC Control Applications of Nonlinear Prog. and Optim.*, Capri, Italy, pp 47–52, 1985.
- [12] Fletcher, R., "A New Approach to Variable Metric Algorithms," *Computer Journal*, Vol. 13, pp 317–322, 1970.
- [13] Fletcher, R., "Practical Methods of Optimization," John Wiley and Sons, 1987.
- [14] Fletcher, R. and M.J.D. Powell, "A Rapidly Convergent Descent Method for Minimization," *Computer Journal*, Vol. 6, pp 163–168, 1963.
- [15] Forsythe, G.F., M.A. Malcolm, and C.B. Moler, *Computer Methods for Mathematical Computations*, Prentice Hall, 1976.
- [16] Gembicki, F.W., "Vector Optimization for Control with Performance and Parameter Sensitivity Indices," Ph.D. Thesis, Case Western Reserve Univ., Cleveland, Ohio, 1974.
- [17] Gill, P.E., W. Murray, M.A. Saunders, and M.H. Wright, "Procedures for Optimization Problems with a Mixture of Bounds and General Linear Constraints," *ACM Trans. Math. Software*, Vol. 10, pp 282–298, 1984.
- [18] Gill, P.E., W. Murray, and M.H. Wright, *Numerical Linear Algebra and Optimization*, Vol. 1, Addison Wesley, 1991.
- [19] Gill, P.E., W. Murray, and M.H. Wright, *Practical Optimization*, London, Academic Press, 1981.
- [20] Goldfarb, D., "A Family of Variable Metric Updates Derived by Variational Means," *Mathematics of Computing*, Vol. 24, pp 23–26, 1970.
- [21] Grace, A.C.W., "Computer-Aided Control System Design Using Optimization Techniques," Ph.D. Thesis, University of Wales, Bangor, Gwynedd, UK, 1989.
- [22] Han, S.P., "A Globally Convergent Method for Nonlinear Programming," *J. Optimization Theory and Applications*, Vol. 22, p. 297, 1977.

- [23] Hock, W. and K. Schittkowski, "A Comparative Performance Evaluation of 27 Nonlinear Programming Codes," *Computing*, Vol. 30, p. 335, 1983.
- [24] Hollingdale, S.H., *Methods of Operational Analysis in Newer Uses of Mathematics* (James Lighthill, ed.), Penguin Books, 1978.
- [25] Levenberg, K., "A Method for the Solution of Certain Problems in Least Squares," *Quart. Appl. Math.* Vol. 2, pp 164–168, 1944.
- [26] Madsen, K. and H. Schjaer-Jacobsen, "Algorithms for Worst Case Tolerance Optimization," *IEEE Transactions of Circuits and Systems*, Vol. CAS-26, Sept. 1979.
- [27] Marquardt, D., "An Algorithm for Least-Squares Estimation of Nonlinear Parameters," *SIAM J. Appl. Math.* Vol. 11, pp 431–441, 1963.
- [28] Moré, J.J., "The Levenberg-Marquardt Algorithm: Implementation and Theory," *Numerical Analysis*, ed. G. A. Watson, Lecture Notes in Mathematics 630, Springer Verlag, pp 105–116, 1977.
- [29] *NAG Fortran Library Manual*, Mark 12, Vol. 4, E04UAF, p. 16.
- [30] Nelder, J.A. and R. Mead, "A Simplex Method for Function Minimization," *Computer J.*, Vol.7, pp 308–313, 1965.
- [31] Nocedal, J. and S. J. Wright. *Numerical Optimization*, Second Edition. Springer Series in Operations Research, Springer Verlag, 2006.
- [32] Powell, M.J.D., "The Convergence of Variable Metric Methods for Nonlinearly Constrained Optimization Calculations," *Nonlinear Programming* 3, (O.L. Mangasarian, R.R. Meyer and S.M. Robinson, eds.), Academic Press, 1978.
- [33] Powell, M.J.D., "A Fast Algorithm for Nonlinearly Constrained Optimization Calculations," *Numerical Analysis*, G.A.Watson ed., Lecture Notes in Mathematics, Springer Verlag, Vol. 630, 1978.
- [34] Powell, M.J.D., "A Fortran Subroutine for Solving Systems of Nonlinear Algebraic Equations," *Numerical Methods for Nonlinear Algebraic Equations*, (P. Rabinowitz, ed.), Ch.7, 1970.

- [35] Powell, M.J.D., "Variable Metric Methods for Constrained Optimization," *Mathematical Programming: The State of the Art*, (A. Bachem, M. Grottschel and B. Korte, eds.) Springer Verlag, pp 288–311, 1983.
- [36] Schittkowski, K., "NLQPL: A FORTRAN-Subroutine Solving Constrained Nonlinear Programming Problems," *Annals of Operations Research*, Vol. 5, pp 485-500, 1985.
- [37] Shanno, D.F., "Conditioning of Quasi-Newton Methods for Function Minimization," *Mathematics of Computing*, Vol. 24, pp 647–656, 1970.
- [38] Waltz, F.M., "An Engineering Approach: Hierarchical Optimization Criteria," *IEEE Trans.*, Vol. AC-12, pp 179–180, April, 1967.
- [39] Branch, M.A., T.F. Coleman, and Y. Li, "A Subspace, Interior, and Conjugate Gradient Method for Large-Scale Bound-Constrained Minimization Problems," *SIAM Journal on Scientific Computing*, Vol. 21, Number 1, pp 1–23, 1999.
- [40] Byrd, R.H., J. C. Gilbert, and J. Nocedal, "A Trust Region Method Based on Interior Point Techniques for Nonlinear Programming," *Mathematical Programming*, Vol 89, No. 1, pp. 149–185, 2000.
- [41] Byrd, R.H., Mary E. Hribar, and Jorge Nocedal, "An Interior Point Algorithm for Large-Scale Nonlinear Programming," *SIAM Journal on Optimization*, Vol 9, No. 4, pp. 877–900, 1999.
- [42] Byrd, R.H., R.B. Schnabel, and G.A. Shultz, "Approximate Solution of the Trust Region Problem by Minimization over Two-Dimensional Subspaces," *Mathematical Programming*, Vol. 40, pp 247–263, 1988.
- [43] Coleman, T.F. and Y. Li, "On the Convergence of Reflective Newton Methods for Large-Scale Nonlinear Minimization Subject to Bounds," *Mathematical Programming*, Vol. 67, Number 2, pp 189–224, 1994.
- [44] Coleman, T.F. and Y. Li, "An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds," *SIAM Journal on Optimization*, Vol. 6, pp 418–445, 1996.

- [45] Coleman, T.F. and Y. Li, “A Reflective Newton Method for Minimizing a Quadratic Function Subject to Bounds on some of the Variables,” *SIAM Journal on Optimization*, Vol. 6, Number 4, pp 1040–1058, 1996.
- [46] Coleman, T.F. and A. Verma, “A Preconditioned Conjugate Gradient Approach to Linear Equality Constrained Minimization,” *Computational Optimization and Applications*, Vol. 20, No. 1, pp. 61–72, 2001.
- [47] Mehrotra, S., “On the Implementation of a Primal-Dual Interior Point Method,” *SIAM Journal on Optimization*, Vol. 2, pp 575–601, 1992.
- [48] Moré, J.J. and D.C. Sorensen, “Computing a Trust Region Step,” *SIAM Journal on Scientific and Statistical Computing*, Vol. 3, pp 553–572, 1983.
- [49] Sorensen, D.C., “Minimization of a Large Scale Quadratic Function Subject to an Ellipsoidal Constraint,” Department of Computational and Applied Mathematics, Rice University, Technical Report TR94-27, 1994.
- [50] Steihaug, T., “The Conjugate Gradient Method and Trust Regions in Large Scale Optimization,” *SIAM Journal on Numerical Analysis*, Vol. 20, pp 626–637, 1983.
- [51] Waltz, R. A. , J. L. Morales, J. Nocedal, and D. Orban, “An interior algorithm for nonlinear optimization that combines line search and trust region steps,” *Mathematical Programming*, Vol 107, No. 3, pp. 391–408, 2006.
- [52] Zhang, Y., “Solving Large-Scale Linear Programs by Interior-Point Methods Under the MATLAB Environment,” Department of Mathematics and Statistics, University of Maryland, Baltimore County, Baltimore, MD, Technical Report TR96-01, July, 1995.
- [53] Hairer, E., S. P. Norsett, and G. Wanner, *Solving Ordinary Differential Equations I - Nonstiff Problems*, Springer-Verlag, pp. 183–184.
- [54] Chvatal, Vasek, *Linear Programming*, W. H. Freeman and Company, 1983.
- [55] Bixby, Robert E., “Implementing the Simplex Method: The Initial Basis,” *ORSA Journal on Computing*, Vol. 4, No. 3, 1992.

- [56] Andersen, Erling D. and Knud D. Andersen, “Presolving in Linear Programming,” *Mathematical Programming*, Vol. 71, pp. 221–245, 1995.
- [57] Lagarias, J. C., J. A. Reeds, M. H. Wright, and P. E. Wright, “Convergence Properties of the Nelder-Mead Simplex Method in Low Dimensions,” *SIAM Journal of Optimization*, Vol. 9, Number 1, pp. 112–147, 1998.
- [58] Dolan, Elizabeth D. , Jorge J. Moré and Todd S. Munson, “Benchmarking Optimization Software with COPS 3.0,” Argonne National Laboratory Technical Report ANL/MCS-TM-273, February 2004.
- [59] Applegate, D. L., R. E. Bixby, V. Chvátal and W. J. Cook, *The Traveling Salesman Problem: A Computational Study*, Princeton University Press, 2007.
- [60] Spellucci, P., “A new technique for inconsistent QP problems in the SQP method,” *Journal of Mathematical Methods of Operations Research*, Volume 47, Number 3, pp. 355–400, October 1998.
- [61] Tone, K., “Revisions of constraint approximations in the successive QP method for nonlinear programming problems,” *Journal of Mathematical Programming*, Volume 26, Number 2, pp. 144–152, June 1983.
- [62] Gondzio, J. “Multiple centrality corrections in a primal dual method for linear programming.” *Computational Optimization and Applications*, Volume 6, Number 2, pp. 137–156, 1996.
- [63] Gould, N. and P. L. Toint. “Preprocessing for quadratic programming.” *Math. Programming*, Series B, Vol. 100, pp. 95–132, 2004.
- [64] Schittkowski, K., “More Test Examples for Nonlinear Programming Codes,” *Lecture Notes in Economics and Mathematical Systems*, Number 282, Springer, p. 45, 1987.

Parallel Computing for Optimization

- “Parallel Computing in Optimization Toolbox Functions” on page 7-2
- “Using Parallel Computing in Optimization Toolbox” on page 7-5
- “Minimizing an Expensive Optimization Problem Using Parallel Computing Toolbox™” on page 7-9
- “Improving Performance with Parallel Computing” on page 7-16

Parallel Computing in Optimization Toolbox Functions

In this section...
“Parallel Optimization Functionality” on page 7-2
“Parallel Estimation of Gradients” on page 7-3
“Nested Parallel Functions” on page 7-4

Parallel Optimization Functionality

Parallel computing is the technique of using multiple processors on a single problem. The reason to use parallel computing is to speed computations.

The Optimization Toolbox solvers `fmincon`, `fgoalattain`, and `fminimax` can automatically distribute the numerical estimation of gradients of objective functions and nonlinear constraint functions to multiple processors. These solvers use parallel gradient estimation under the following conditions:

- You have a license for Parallel Computing Toolbox software.
- The option `GradObj` is set to `'off'`, or, if there is a nonlinear constraint function, the option `GradConstr` is set to `'off'`. Since `'off'` is the default value of these options, you don't have to set them; just don't set them both to `'on'`.
- Parallel computing is enabled with `parpool`, a Parallel Computing Toolbox function.
- The option `UseParallel` is set to `true`. The default value of this option is `false`.

When these conditions hold, the solvers compute estimated gradients in parallel.

Note Even when running in parallel, a solver occasionally calls the objective and nonlinear constraint functions serially on the host machine. Therefore, ensure that your functions have no assumptions about whether they are evaluated in serial or parallel.

Parallel Estimation of Gradients

One subroutine was made parallel in the functions `fmincon`, `fgoalattain`, and `fminimax`: the subroutine that estimates the gradient of the objective function and constraint functions. This calculation involves computing function values at points near the current location x . Essentially, the calculation is

$$\nabla f(x) \approx \left[\frac{f(x + \Delta_1 e_1) - f(x)}{\Delta_1}, \frac{f(x + \Delta_2 e_2) - f(x)}{\Delta_2}, \dots, \frac{f(x + \Delta_n e_n) - f(x)}{\Delta_n} \right],$$

where

- f represents objective or constraint functions
- e_i are the unit direction vectors
- Δ_i is the size of a step in the e_i direction

To estimate $\nabla f(x)$ in parallel, Optimization Toolbox solvers distribute the evaluation of $(f(x + \Delta_i e_i) - f(x))/\Delta_i$ to extra processors.

Parallel Central Differences

You can choose to have gradients estimated by central finite differences instead of the default forward finite differences. The basic central finite difference formula is

$$\nabla f(x) \approx \left[\frac{f(x + \Delta_1 e_1) - f(x - \Delta_1 e_1)}{2\Delta_1}, \dots, \frac{f(x + \Delta_n e_n) - f(x - \Delta_n e_n)}{2\Delta_n} \right].$$

This takes twice as many function evaluations as forward finite differences, but is usually much more accurate. Central finite differences work in parallel exactly the same as forward finite differences.

Enable central finite differences by using `optimoptions` to set the `FinDiffType` option to 'central'. To use forward finite differences, set the `FinDiffType` option to 'forward'.

Nested Parallel Functions

Solvers employ the Parallel Computing Toolbox function `parfor` to perform parallel estimation of gradients. `parfor` does not work in parallel when called from within another `parfor` loop. Therefore, you cannot simultaneously use parallel gradient estimation and parallel functionality within your objective or constraint functions.

Suppose, for example, your objective function `userfcn` calls `parfor`, and you wish to call `fmincon` in a loop. Suppose also that the conditions for parallel gradient evaluation of `fmincon`, as given in “Parallel Optimization Functionality” on page 7-2, are satisfied. When `parfor` Runs In Parallel on page 7-4 shows three cases:

- 1 The outermost loop is `parfor`. Only that loop runs in parallel.
- 2 The outermost `parfor` loop is in `fmincon`. Only `fmincon` runs in parallel.
- 3 The outermost `parfor` loop is in `userfcn`. `userfcn` can use `parfor` in parallel.

Bold indicates the function that runs in parallel

- ① `... parfor (i=1:10)`
`x(i)=fmincon(@userfcn,...)`
`...`
`end`
 Only the outermost `parfor` loop runs in parallel
- ② `... for (i=1:10)`
`x(i)=fmincon(@userfcn,...)`
`...`
`end`
 If `UseParallel` = 'always' `fmincon` runs in parallel
- ③ `... for (i=1:10)`
`x(i)=fmincon(@userfcn,...)`
`...`
`end`
 If `UseParallel` = 'never' `userfcn` can use `parfor` in parallel

When `parfor` Runs In Parallel

Using Parallel Computing in Optimization Toolbox

In this section...

“Using Parallel Computing with Multicore Processors” on page 7-5

“Using Parallel Computing with a Multiprocessor Network” on page 7-6

“Testing Parallel Computations” on page 7-7

Using Parallel Computing with Multicore Processors

If you have a multicore processor, you might see speedup using parallel processing. You can establish a parallel pool of several workers with a Parallel Computing Toolbox license. For a description of Parallel Computing Toolbox software, see “Getting Started with Parallel Computing Toolbox”.

Suppose you have a dual-core processor, and want to use parallel computing:

- Enter

```
parpool
```

at the command line. MATLAB starts a pool of workers using the multicore processor. If you had previously set a nondefault cluster profile, you can enforce multicore (local) computing:

```
parpool('local')
```

Note Depending on your preferences, MATLAB can start a parallel pool automatically. To enable this feature, check **Automatically create a parallel pool** in **Home > Parallel > Parallel Preferences**.

- - For command-line use, enter


```
options = optimoptions('solvername','UseParallel',true);
```
 - For Optimization app, check **Options > Approximated derivatives > Evaluate in parallel**.

When you run an applicable solver with options, applicable solvers automatically use parallel computing.

To stop computing optimizations in parallel, set `UseParallel` to `false`, or set the Optimization app not to compute in parallel. To halt all parallel computation, enter

```
delete(gcp)
```

Using Parallel Computing with a Multiprocessor Network

If you have multiple processors on a network, use Parallel Computing Toolbox functions and MATLAB Distributed Computing Server™ software to establish parallel computation. Here are the steps to take:

- 1** Make sure your system is configured properly for parallel computing. Check with your systems administrator, or refer to the Parallel Computing Toolbox documentation.

To perform a basic check:

- a** At the command line, enter

```
parpool('prof')
```

where `prof` is your cluster profile.

- b** Workers must be able to access your objective function file and, if applicable, your nonlinear constraint function file. There are two ways of ensuring access:
 - i** Distribute the files to the workers using the `parpool` `AttachedFiles` argument. For example, if `objfun.m` is your objective function file, and `constrfun.m` is your nonlinear constraint function file, enter

```
parpool('AttachedFiles',{ 'objfun.m', 'constrfun.m' });
```

Workers access their own copies of the files.

- ii** Give a network file path to your files. If `network_file_path` is the network path to your objective or constraint function files, enter

```
pctRunOnAll('addpath network_file_path')
```


Workers access the function files over the network.

- c Check whether a file is on the path of every worker by entering

```
pctRunOnAll('which filename')
```

If any worker does not have a path to the file, it reports

```
filename not found.
```

2

- For command-line use, enter

```
options = optimoptions('solvername','UseParallel',true);
```

- For Optimization app, check **Options > Approximated derivatives > Evaluate in parallel**.

After you establish your parallel computing environment, applicable solvers automatically use parallel computing whenever you call them with options.

To stop computing optimizations in parallel, set `UseParallel` to `false`, or set the Optimization app not to compute in parallel. To halt all parallel computation, enter

```
delete(gcp)
```

Testing Parallel Computations

To test see if a problem runs correctly in parallel,

- 1 Try your problem without parallel computation to ensure that it runs properly serially. Make sure this is successful (gives correct results) before going to the next test.
- 2 Set `UseParallel` to `true`, and ensure that there is no parallel pool using `delete(gcp)`. Uncheck **Automatically create a parallel pool** in **Home > Parallel > Parallel Preferences** so MATLAB does not create a parallel pool. Your problem runs `parfor` serially, with loop iterations in reverse order from a `for` loop. Make sure this is successful (gives correct results) before going to the next test.

- 3** Set `UseParallel` to `true`, and create a parallel pool using `parpool`. Unless you have a multicore processor or a network set up, you won't see any speedup. This testing is simply to verify the correctness of the computations.

Remember to call your solver using an options structure to test or use parallel functionality.

Minimizing an Expensive Optimization Problem Using Parallel Computing Toolbox™

This example shows how to speed up the minimization of an expensive optimization problem using functions in Optimization Toolbox™ and Global Optimization Toolbox. In the first part of the example we solve the optimization problem by evaluating functions in a serial fashion and in the second part of the example we solve the same problem using the parallel for loop (parfor) feature by evaluating functions in parallel. We compare the time taken by the optimization function in both cases.

Expensive Optimization Problem

For the purpose of this example, we solve a problem in four variables, where the objective and constraint functions are made artificially expensive by pausing.

```
type expensive_objfun.m
type expensive_confun.m
```

```
function f = expensive_objfun(x)
%EXPENSIVE_OBFUN An expensive objective function used in optimparfor exampl

% Copyright 2007-2013 The MathWorks, Inc.
% $Revision: 1.1.8.2 $ $Date: 2013/05/04 00:47:14 $

% Simulate an expensive function by pausing
pause(0.1)
% Evaluate objective function
f = exp(x(1)) * (4*x(3)^2 + 2*x(4)^2 + 4*x(1)*x(2) + 2*x(2) + 1);

function [c,ceq] = expensive_confun(x)
%EXPENSIVE_CONFUN An expensive constraint function used in optimparfor exampl

% Copyright 2007-2013 The MathWorks, Inc.
% $Revision: 1.1.8.2 $ $Date: 2013/05/04 00:47:13 $

% Simulate an expensive function by pausing
pause(0.1);
```

```
% Evaluate constraints
c = [1.5 + x(1)*x(2)*x(3) - x(1) - x(2) - x(4);
     -x(1)*x(2) + x(4) - 10];
% No nonlinear equality constraints:
ceq = [];
```

Minimizing Using fmincon

We are interested in measuring the time taken by fmincon in serial so that we can compare it to the parallel fmincon evaluation.

```
startPoint = [-1 1 1 -1];
options = optimoptions('fmincon','Display','iter','Algorithm','sqp');
startTime = tic;
xsol = fmincon(@expensive_objfun,startPoint,[],[],[],[],[],[],[],@expensive_co
time_fmincon_sequential = toc(startTime);
fprintf('Serial FMINCON optimization takes %g seconds.\n',time_fmincon_sequ
```

Iter	F-count	f(x)	Feasibility	Steplength	Norm of First-order step	optimality
0	5	1.839397e+00	1.500e+00			3.311e+
1	12	-8.841073e-01	4.019e+00	4.900e-01	2.335e+00	7.015e-
2	17	-1.382832e+00	0.000e+00	1.000e+00	1.142e+00	9.272e-
3	22	-2.241952e+00	0.000e+00	1.000e+00	2.447e+00	1.481e+
4	27	-3.145762e+00	0.000e+00	1.000e+00	1.756e+00	5.464e+
5	32	-5.277523e+00	6.413e+00	1.000e+00	2.224e+00	1.357e+
6	37	-6.310709e+00	0.000e+00	1.000e+00	1.099e+00	1.309e+
7	43	-6.447956e+00	0.000e+00	7.000e-01	2.191e+00	3.631e+
8	48	-7.135133e+00	0.000e+00	1.000e+00	3.719e-01	1.205e-
9	53	-7.162732e+00	0.000e+00	1.000e+00	4.083e-01	2.935e-
10	58	-7.178390e+00	0.000e+00	1.000e+00	1.591e-01	3.110e-
11	63	-7.180399e+00	1.191e-05	1.000e+00	2.644e-02	1.553e-
12	68	-7.180408e+00	0.000e+00	1.000e+00	1.140e-02	5.584e-
13	73	-7.180411e+00	0.000e+00	1.000e+00	1.764e-03	4.677e-
14	78	-7.180412e+00	0.000e+00	1.000e+00	8.827e-05	1.304e-
15	83	-7.180412e+00	0.000e+00	1.000e+00	1.528e-06	1.023e-

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in

feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint

Serial FMINCON optimization takes 18.1397 seconds.

Minimizing Using Genetic Algorithm

Since `ga` usually takes many more function evaluations than `fmincon`, we remove the expensive constraint from this problem and perform unconstrained optimization instead; we pass empty (`[]`) for constraints. In addition, we limit the maximum number of generations to 15 for `ga` so that `ga` can terminate in a reasonable amount of time. We are interested in measuring the time taken by `ga` so that we can compare it to the parallel `ga` evaluation. Note that running `ga` requires Global Optimization Toolbox.

```
rng default % for reproducibility
try
    gaAvailable = false;
    nvar = 4;
    gaoptions = gaoptimset('Generations',15,'Display','iter');
    startTime = tic;
    gasol = ga(@expensive_objfun,nvar,[],[],[],[],[],[],[],[],gaoptions);
    time_ga_sequential = toc(startTime);
    fprintf('Serial GA optimization takes %g seconds.\n',time_ga_sequential);
    gaAvailable = true;
catch ME
    warning(message('optimdemos:optimparfor:gaNotFound'));
end
```

Generation	f-count	Best f(x)	Mean f(x)	Stall Generations
1	100	-6.433e+16	-1.287e+15	0
2	150	-1.501e+17	-7.138e+15	0
3	200	-7.878e+26	-1.576e+25	0
4	250	-8.664e+27	-1.466e+26	0
5	300	-1.096e+28	-2.062e+26	0
6	350	-5.422e+33	-1.145e+32	0

7	400	-1.636e+36	-3.316e+34	0
8	450	-2.933e+36	-1.513e+35	0
9	500	-1.351e+40	-2.705e+38	0
10	550	-1.351e+40	-7.9e+38	1
11	600	-2.07e+40	-2.266e+39	0
12	650	-1.845e+44	-3.696e+42	0
13	700	-2.893e+44	-1.687e+43	0
14	750	-5.076e+44	-6.516e+43	0
15	800	-8.321e+44	-2.225e+44	0

Optimization terminated: maximum number of generations exceeded.
Serial GA optimization takes 87.3686 seconds.

Setting Parallel Computing Toolbox

The finite differencing used by the functions in Optimization Toolbox to approximate derivatives is done in parallel using the `parfor` feature if Parallel Computing Toolbox is available and there is a parallel pool of workers. Similarly, `ga`, `gamultiobj`, and `patternsearch` solvers in Global Optimization Toolbox evaluate functions in parallel. To use the `parfor` feature, we use the `parpool` function to set up the parallel environment. The computer on which this example is published has four cores, so `parpool` starts four MATLAB® workers. If there is already a parallel pool when you run this example, we use that pool; see the documentation for `parpool` for more information.

```
if max(size(gcf)) == 0 % parallel pool needed
    parpool % create the parallel pool
end
```

```
Starting parallel pool (parpool) using the 'local' profile ... connected to
```

Minimizing Using Parallel `fmincon`

To minimize our expensive optimization problem using the parallel `fmincon` function, we need to explicitly indicate that our objective and constraint functions can be evaluated in parallel and that we want `fmincon` to use its parallel functionality wherever possible. Currently, finite differencing can be done in parallel. We are interested in measuring the time taken by `fmincon` so that we can compare it to the serial `fmincon` run.

```
options = optimoptions(options,'UseParallel',true);
```

```

startTime = tic;
xsol = fmincon(@expensive_objfun,startPoint,[],[],[],[],[],[],[],@expensive_co
time_fmincon_parallel = toc(startTime);
fprintf('Parallel FMINCON optimization takes %g seconds.\n',time_fmincon_pa

```

Iter	F-count	f(x)	Feasibility	Steplength	Norm of step	First-order optimality
0	5	1.839397e+00	1.500e+00			3.311e+
1	12	-8.841073e-01	4.019e+00	4.900e-01	2.335e+00	7.015e-
2	17	-1.382832e+00	0.000e+00	1.000e+00	1.142e+00	9.272e-
3	22	-2.241952e+00	0.000e+00	1.000e+00	2.447e+00	1.481e+
4	27	-3.145762e+00	0.000e+00	1.000e+00	1.756e+00	5.464e+
5	32	-5.277523e+00	6.413e+00	1.000e+00	2.224e+00	1.357e+
6	37	-6.310709e+00	0.000e+00	1.000e+00	1.099e+00	1.309e+
7	43	-6.447956e+00	0.000e+00	7.000e-01	2.191e+00	3.631e+
8	48	-7.135133e+00	0.000e+00	1.000e+00	3.719e-01	1.205e-
9	53	-7.162732e+00	0.000e+00	1.000e+00	4.083e-01	2.935e-
10	58	-7.178390e+00	0.000e+00	1.000e+00	1.591e-01	3.110e-
11	63	-7.180399e+00	1.191e-05	1.000e+00	2.644e-02	1.553e-
12	68	-7.180408e+00	0.000e+00	1.000e+00	1.140e-02	5.584e-
13	73	-7.180411e+00	0.000e+00	1.000e+00	1.764e-03	4.677e-
14	78	-7.180412e+00	0.000e+00	1.000e+00	8.827e-05	1.304e-
15	83	-7.180412e+00	0.000e+00	1.000e+00	1.528e-06	1.023e-

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint

Parallel FMINCON optimization takes 8.78988 seconds.

Minimizing Using Parallel Genetic Algorithm

To minimize our expensive optimization problem using the `ga` function, we need to explicitly indicate that our objective function can be evaluated in parallel and that we want `ga` to use its parallel functionality wherever possible. To use the parallel `ga` we also require that the 'Vectorized' option

be set to the default (i.e., 'off'). We are again interested in measuring the time taken by `ga` so that we can compare it to the serial `ga` run. Though this run may be different from the serial one because `ga` uses a random number generator, the number of expensive function evaluations is the same in both runs. Note that running `ga` requires Global Optimization Toolbox.

```
rng default % to get the same evaluations as the previous run
if gaAvailable
    gaoptions = gaoptimset(gaoptions,'UseParallel',true);
    startTime = tic;
    gasol = ga(@expensive_objfun,nvar,[],[],[],[],[],[],[],gaoptions);
    time_ga_parallel = toc(startTime);
    fprintf('Parallel GA optimization takes %g seconds.\n',time_ga_parallel)
end
```

Generation	f-count	Best f(x)	Mean f(x)	Stall Generations
1	100	-6.433e+16	-1.287e+15	0
2	150	-1.501e+17	-7.138e+15	0
3	200	-7.878e+26	-1.576e+25	0
4	250	-8.664e+27	-1.466e+26	0
5	300	-1.096e+28	-2.062e+26	0
6	350	-5.422e+33	-1.145e+32	0
7	400	-1.636e+36	-3.316e+34	0
8	450	-2.933e+36	-1.513e+35	0
9	500	-1.351e+40	-2.705e+38	0
10	550	-1.351e+40	-7.9e+38	1
11	600	-2.07e+40	-2.266e+39	0
12	650	-1.845e+44	-3.696e+42	0
13	700	-2.893e+44	-1.687e+43	0
14	750	-5.076e+44	-6.516e+43	0
15	800	-8.321e+44	-2.225e+44	0

Optimization terminated: maximum number of generations exceeded.
Parallel GA optimization takes 23.707 seconds.

Compare Serial and Parallel Time

```
X = [time_fmincon_sequential time_fmincon_parallel];
Y = [time_ga_sequential time_ga_parallel];
```



```
t = [0 1];
plot(t,X,'r--',t,Y,'k-')
ylabel('Time in seconds')
legend('fmincon','ga')
set(gca,'XTick',[0 1])
set(gca,'XTickLabel',{'Serial' 'Parallel'})
axis([0 1 0 ceil(max([X Y]))])
title('Serial Vs. Parallel Times')
```

Utilizing parallel function evaluation via `parfor` improved the efficiency of both `fmincon` and `ga`. The improvement is typically better for expensive objective and constraint functions.

At last we delete the parallel pool.

```
if max(size(gcf)) > 0 % parallel pool exists
    delete(gcf) % delete the pool
end
```

Parallel pool using the 'local' profile is shutting down.

Improving Performance with Parallel Computing

In this section...

“Factors That Affect Speed” on page 7-16

“Factors That Affect Results” on page 7-17

“Searching for Global Optima” on page 7-17

Factors That Affect Speed

Some factors may affect the speed of execution of parallel processing:

- **Parallel overhead.** There is overhead in calling `parfor` instead of `for`. If function evaluations are fast, this overhead could become appreciable. In particular, solving a problem in parallel can be slower than solving the problem serially.
- **No nested `parfor` loops.** This is described in “Nested Parallel Functions” on page 7-4. `parfor` does not work in parallel when called from within another `parfor` loop. If you have programmed your objective or constraint functions to take advantage of parallel processing, the limitation of no nested `parfor` loops may cause a solver to run more slowly than you expect. In particular, the parallel computation of finite differences takes precedence, since that is an outer loop. This causes any parallel code within the objective or constraint functions to execute serially.
- **When executing serially, `parfor` loops run slower than `for` loops.** Therefore, for best performance, ensure that only your outermost parallel loop calls `parfor`. For example, suppose your code calls `fmincon` within a `parfor` loop. For best performance in this case, set the `fmincon` `UseParallel` option to `false`.
- **Passing parameters.** Parameters are automatically passed to worker machines during the execution of parallel computations. If there are a large number of parameters, or they take a large amount of memory, passing them may slow the execution of your computation.
- **Contention for resources: network and computing.** If the network of worker machines has low bandwidth or high latency, computation could be slowed.

Factors That Affect Results

Some factors may affect numerical results when using parallel processing. There are more caveats related to `parfor` listed in the “Limitations” section of the Parallel Computing Toolbox documentation.

- Persistent or global variables. If your objective or constraint functions use persistent or global variables, these variables may take different values on different worker processors. Furthermore, they may not be cleared properly on the worker processors.
- Accessing external files. External files may be accessed in an unpredictable fashion during a parallel computation. The order of computations is not guaranteed during parallel processing, so external files may be accessed in unpredictable order, leading to unpredictable results.
- Accessing external files. If two or more processors try to read an external file simultaneously, the file may become locked, leading to a read error, and halting the execution of the optimization.
- If your objective function calls Simulink, results may be unreliable with parallel gradient estimation.
- Noncomputational functions, such as `input`, `plot`, and `keyboard`, might behave badly when used in objective or constraint functions. When called in a `parfor` loop, these functions are executed on worker machines. This can cause a worker to become nonresponsive, since it is waiting for input.
- `parfor` does not allow `break` or `return` statements.

Searching for Global Optima

To search for global optima, one approach is to evaluate a solver from a variety of initial points. If you distribute those evaluations over a number of processors using the `parfor` function, you disable parallel gradient estimation, since `parfor` loops cannot be nested. Your optimization usually runs more quickly if you distribute the evaluations over all the processors, rather than running them serially with parallel gradient estimation, so disabling parallel estimation probably won't slow your computation. If you have more processors than initial points, though, it is not clear whether it is better to distribute initial points or to enable parallel gradient estimation.

If you have a Global Optimization Toolbox license, you can use the `MultiStart` solver to examine multiple start points in parallel. See “Parallel Computing” and “Parallel MultiStart” in the Global Optimization Toolbox documentation.

External Interface

ktrlink: An Interface to KNITRO Libraries

In this section...

- “What Is ktrlink?” on page 8-2
- “Installation and Configuration” on page 8-2
- “Using ktrlink to Optimize” on page 8-4
- “Setting Options” on page 8-8
- “Sparse Matrix Considerations” on page 8-9

What Is ktrlink?

ktrlink calls Ziena Optimization’s KNITRO® libraries in order to perform an optimization. ktrlink can address constrained and unconstrained problems. To use ktrlink, you must purchase a copy of KNITRO libraries from Ziena Optimization, Inc. (<http://www.ziena.com/>).

Use ktrlink the same as any other Optimization Toolbox function.

ktrlink’s syntax is similar to fmincon’s. The main differences are:

- ktrlink has additional options input for KNITRO libraries so you can access its options.
- ktrlink has no provision for obtaining a returned Hessian or gradient, since KNITRO software doesn’t return them.
- Sparse matrix representations differ between KNITRO software and MATLAB.

Furthermore, many returned flags and messages differ from fmincon’s, because they are returned directly from KNITRO libraries.

Installation and Configuration

The system requirements for MATLAB and KNITRO software may differ. Check the system requirements for both products before attempting to use ktrlink. For recent and planned MATLAB platform support, see

<http://www.mathworks.com/support/sysreq/roadmap.html>

Optimization Toolbox software version 6 works with versions 6, 7, and 8 of KNITRO libraries. Contact Ziena Optimization, Inc. (<http://www.ziena.com/>) if you have questions regarding other versions of the KNITRO libraries.

Perform the following steps to configure your system to use ktrlink:

- 1 Install MATLAB and the KNITRO libraries on your system.
- 2 Set the system path to include the KNITRO libraries (see “Setting the System Path to Include KNITRO Libraries” on page 8-3). Make sure to perform this step before starting MATLAB.
- 3 Start MATLAB.

Setting the System Path to Include KNITRO Libraries

In order to use ktrlink, you need to tell MATLAB where the KNITRO binary file (`libknitro.so`, `libknitro.dylib`, `knitro.dll`, or a similar file) resides. You do this by setting a system-wide environment variable. Enter the following system-level commands. Replace `<file_absolute_path>` with the full path to your KNITRO libraries:

- Linux:

```
setenv LD_LIBRARY_PATH <file_absolute_path>:$LD_LIBRARY_PATH
```

- Macintosh:

- 1 Launch **Terminal**, available in **Applications > Utilities**.

- 2 Determine the shell:

```
echo $SHELL
```

- 3 Set the path:

- If shell is bash:

```
export DYLD_LIBRARY_PATH=<file_absolute_path>:$DYLD_LIBRARY_PATH
```

- If shell is `csh` or `tcsh`:

```
setenv DYLD_LIBRARY_PATH=<file_absolute_path>:$DYLD_LIBRARY_PATH
```

- 4 Launch MATLAB by entering `matlab` in **Terminal**. (Launching MATLAB by clicking an icon can lose the path to the KNITRO binary file.)

Optional: to set the path permanently:

- In `bash`, add the `export` line in step 3 to `$HOME/.profile`.
 - In `csh` or `tcsh`, add the `setenv` line in step 3 to `$HOME/.cshrc`.
- Windows:
 - 1 At the Windows desktop, right-click **My Computer** (Windows XP) or **Computer** (Vista or Windows 7).
 - 2 Select **Properties**.
 - 3 Click the **Advanced** tab (Windows XP) or **Advanced System Settings** (Vista or Windows 7).
 - 4 Click **Environment Variables**.
 - 5 Under **System variables**, edit the **Path** variable to add the KNITRO library folder.

Check if the installation was successful by starting MATLAB and running the following command:

```
[x fval] = ktrlink(@(x)cos(x),1)
```

If you receive an error message, check your system path, and make sure the KNITRO libraries are on the path. When installed correctly, `ktrlink` returns a long message, ending with `x = 3.1416`, `fval = -1`.

Using `ktrlink` to Optimize

- 1 This example uses the same constraint function as the example in “Nonlinear Constraints” on page 2-37. The constraint function is the intersection of the interior of an ellipse with the region above a parabola:

```
function [c ceq gradc gradceq]=ellipseparabola(x)
```



```

% Inside the ellipse bounded by (-3<x<3),(-2<y<2)
% Above the line y=x^2-1
c(1) = x(1)^2/9 + x(2)^2/4 - 1;% ellipse
c(2) = x(1)^2 - x(2) - 1;% parabola
ceq = [];

if nargout > 2
    gradc = [2*x(1)/9, 2*x(1);...
            x(2)/2, -1];
    gradceq = [];
end

```

- 2** The objective function is a tilted sinh:

```

function [f gradf]=sinhtilt(x)

A=[2,1;1,2];
m=[1,1];
f=sinh(x'*A*x/100) + sinh(m*A*x/10);

if nargout > 1
    gradf=cosh(x'*A*x/100)*(A*x)/50;
    gradf=gradf+cosh(m*A*x/10)*[3;3]/10;
end

```

- 3** Set the options so that ktrlink has iterative display and uses the gradients included in the objective and constraint functions:

```

kthropts = optimoptions('ktrlink','Display','iter',...
    'GradConstr','on','GradObj','on');

```

- 4** Run the optimization starting at [0;0], using the kthropts options:

```

[x fval flag] = ktrlink(@sinhtilt,[0;0],...
    [],[],[],[],[],[],@ellipseparabola,kthropts)

```

KNITRO software returns the following output:

```

=====
Trial Zienna License (NOT FOR COMMERCIAL USE)
      KNITRO 8.1.0
      Zienna Optimization

```

=====

KNITRO presolve eliminated 0 variables and 0 constraints.

```

algorithm:          1
hessopt:           2
honorbnds:         1
maxit:             10000
outlev:            4
KNITRO changing bar_switchrule from AUTO to 2.
KNITRO changing bar_murule from AUTO to 4.
KNITRO changing bar_initpt from AUTO to 2.
KNITRO changing bar_penaltyrule from AUTO to 2.
KNITRO changing bar_penaltycons from AUTO to 1.
KNITRO changing linsolver from AUTO to 2.
    
```

Problem Characteristics

```

Objective goal: Minimize
Number of variables:          2
    bounded below:           0
    bounded above:           0
    bounded below and above:  0
    fixed:                   0
    free:                     2
Number of constraints:        2
    linear equalities:        0
    nonlinear equalities:     0
    linear inequalities:      0
    nonlinear inequalities:   2
    range:                   0
Number of nonzeros in Jacobian: 4
Number of nonzeros in Hessian: 3
    
```

Iter	fCount	Objective	FeasError	OptError	Step	CGits
0	1	0.000000e+000	0.000e+000			
1	2	-1.701203e-001	0.000e+000	1.969e-001	4.111e-001	0
2	3	-4.135676e-001	2.440e-001	1.356e-001	6.054e-001	0
3	4	-3.631949e-001	2.037e-002	4.434e-002	1.423e-001	0

4	5	-3.608513e-001	3.776e-003	1.680e-002	8.094e-002	0
5	6	-3.604610e-001	1.216e-003	1.494e-003	4.870e-002	0
6	7	-3.601203e-001	2.694e-006	1.630e-004	3.329e-003	0
7	8	-3.601196e-001	0.000e+000	4.840e-007	1.241e-005	0

EXIT: Locally optimal solution found.

Final Statistics

Final objective value = -3.60119566305939e-001
Final feasibility error (abs / rel) = 0.00e+000 / 0.00e+000
Final optimality error (abs / rel) = 4.84e-007 / 4.84e-007
of iterations = 7
of CG iterations = 0
of function evaluations = 8
of gradient evaluations = 8
Total program time (secs) = 0.097 (0.218 CPU time)
Time spent in evaluations (secs) = 0.067

=====

x =

-0.5083
-0.7416

fval =

-0.3601

flag =

0

Note Exit flags have different meanings for ktrlink and fmincon. Flag 0 for KNITRO libraries means the first-order optimality condition was satisfied; for fmincon, the corresponding flag is 1. For more information about the output, see the KNITRO documentation at <http://www.ziena.com/>.

Setting Options

`ktrlink` takes up to two options inputs: one in `fmincon` format, and another in KNITRO format. You can use either or both types of options. If you use both types of options, MATLAB reads only four `fmincon`-format options: `HessFcn`, `HessMult`, `HessPattern`, and `JacobPattern`. KNITRO options override `fmincon`-format options.

To use KNITRO options, create an options text file, whose format can be found in the KNITRO documentation. For example, if you have a KNITRO options file named `knitropts`, and an `fmincon`-format options structure named `ktropts`, you can pass them both by calling `ktrlink` like this:

```
[x fval] = ktrlink(@bigtopleft,[-1,-1,-1],...
                 [],[],[],[],[],[],@twocone,ktropts,'knitropts')
```

If `knitropts` resides in a different folder, pass the path to the file. For example:

```
[x fval] = ktrlink(@bigtopleft,[-1,-1,-1],...
                 [],[],[],[],[],[],@twocone,ktropts,...
                 'C:\Documents\Knitro\knitropts')
```

The following shows how `fmincon`-format options are mapped to KNITRO options.

fmincon Option	KNITRO Option	Default Value
Algorithm	algorithm	'interior-point'
AlwaysHonorConstraints	honorbounds	'bounds'
Display	outlev	'none'
FinDiffType	gradopt	'forward'
GradConstr	gradopt	'off'
GradObj	gradopt	'off'
HessFcn	hessopt	[]
MaxIter	maxit	10000
TolFun	opttol	1.00E-06
TolX	xtol	1.00E-15

fmincon Option	KNITRO Option	Default Value
Hessian, LBFGS pairs	hessopt, lmsize	'bfgs'
HessMult	hessopt	[]
HessPattern		[]
InitBarrierParam	bar_initmu	0.1
InitTrustRegionRadius	delta	sqrt(numberOfVariables)
JacobPattern		[]
MaxProjCGIter	maxcgit	2*(numberOfVariables–numberOfEqualities)
ObjectiveLimit	objrange	-1.0E20
ScaleProblem	scale	'obj-and-constr'
SubProblemAlgorithm	algorithm	'ldl-factorization'
TolCon	feastol	1.00E-06

Note KNITRO libraries allow you to pass simultaneously either the gradients of the objective function and all nonlinear constraints, or no gradients. Therefore, setting `GradObj = 'on'` and `GradConstr = 'off'` is inconsistent. If you attempt to pass inconsistent options, `ktrlink` warns you, and treats all gradients as if they had not been passed.

Sparse Matrix Considerations

When the Hessian of the Lagrangian is sparse, or the Jacobian of the nonlinear constraints is sparse, `ktrlink` makes use of the sparsity structure to speed the optimization and use less memory doing so.

`ktrlink` handles sparse matrices differently than other MATLAB functions. If you choose to use sparse matrices, `ktrlink` requires a sparsity pattern for nonlinear constraint Jacobians and for Hessians. The next two sections give the details of the sparsity patterns for nonlinear constraint Jacobians and for Hessians.

Sparsity Pattern for Nonlinear Constraints

The sparsity pattern for constraint Jacobians is a matrix. You pass the matrix as the `JacobPattern` option. The structure of the matrix follows.

Let c denote the vector of m nonlinear inequality constraints, and let ceq denote the vector of m_2 nonlinear equality constraints. If there are n dimensions, the Jacobian is an $(m + m_2)$ -by- n matrix. The Jacobian pattern J is

$$J_{i,j} = \begin{cases} 1 & \text{if } \frac{\partial c_i(x)}{\partial x_j} \neq 0 \\ 0 & \text{if } \frac{\partial c_i(x)}{\partial x_j} \equiv 0 \end{cases}, \quad 1 \leq i \leq m,$$

$$J_{i+m,j} = \begin{cases} 1 & \text{if } \frac{\partial ceq_i(x)}{\partial x_j} \neq 0 \\ 0 & \text{if } \frac{\partial ceq_i(x)}{\partial x_j} \equiv 0 \end{cases}, \quad 1 \leq i \leq m_2.$$

In other words, the i th row of the Jacobian pattern corresponds to the gradient of c_i . Inequality gradients lie above equality gradients (they have lower row numbers).

All that matters for the Jacobian pattern is whether or not the entries are zero. You can pass single-precision numbers, doubles, or logical true or false. You can pass the sparsity pattern as a MATLAB sparse matrix. If you have a large sparse matrix of constraints, it is more efficient to pass the pattern as a sparse matrix. Linear constraint matrices are automatically passed as sparse matrices.

The gradient of the constraints, calculated in the constraint function, has the transpose of the Jacobian pattern. For more information about the form of constraint gradients, see “Nonlinear Constraints” on page 2-37.

Sparsity Pattern for Hessians

The Hessian is the matrix of second derivatives of the Lagrangian:

$$H_{i,j} = \frac{\partial^2 f}{\partial x_i \partial x_j} + \sum_k \lambda_k \frac{\partial^2 c_k}{\partial x_i \partial x_j} + \sum_m \lambda_m \frac{\partial^2 ceq_m}{\partial x_i \partial x_j}.$$

Give the matrix H as a full or sparse matrix of zero and nonzero elements. The elements can be single-precision numbers, doubles, or logical true or false.

The Hessian is a symmetric matrix. You can pass just the upper triangular pattern, or pass the whole matrix pattern.

Argument and Options Reference

- “Function Arguments” on page 9-2
- “Optimization Options Reference” on page 9-7

Function Arguments

In this section...
“Input Arguments” on page 9-2
“Output Arguments” on page 9-5

Input Arguments

Argument	Description	Used by Functions
A, b	The matrix A and vector b are, respectively, the coefficients of the linear inequality constraints and the corresponding right-side vector: $A*x \leq b$.	bintprog, fgoalattain, fmincon, fminimax, fsemif, linprog, lsqlin, quadprog
Aeq, beq	The matrix Aeq and vector beq are, respectively, the coefficients of the linear equality constraints and the corresponding right-side vector: $Aeq*x = beq$.	bintprog, fgoalattain, fmincon, fminimax, fsemif, linprog, lsqlin, quadprog
C, d	The matrix C and vector d are, respectively, the coefficients of the over or underdetermined linear system and the right-side vector to be solved.	lsqlin, lsqnonneg
f	The vector of coefficients for the linear term in the linear equation $f' * x$ or the quadratic equation $x' * H * x + f' * x$.	bintprog, linprog, quadprog
fun	The function to be optimized. fun is either a function handle to a file or is an anonymous function. See the individual function reference pages for more information on fun.	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fsemif, fsolve, fzero, lsqcurvefit, lsqnonlin

Argument	Description	Used by Functions
goal	Vector of values that the objectives attempt to attain. The vector is the same length as the number of objectives.	fgoalattain
H	The matrix of coefficients for the quadratic terms in the quadratic equation $x' * H * x + f' * x$. H must be symmetric.	quadprog
lb, ub	Lower and upper bound vectors (or matrices). The arguments are normally the same size as x. However, if lb has fewer elements than x, say only m, then only the first m elements in x are bounded below; upper bounds in ub can be defined in the same manner. You can also specify unbounded variables using -Inf (for lower bounds) or Inf (for upper bounds). For example, if $lb(i) = -Inf$, the variable $x(i)$ is unbounded below.	fgoalattain, fmincon, fminimax, fseminf, linprog, lsqcurvefit, lsqlin, lsqnonlin, quadprog
nonlcon	The function that computes the nonlinear inequality and equality constraints. "Passing Extra Parameters" on page 2-53 explains how to parameterize the function nonlcon, if necessary. See the individual reference pages for more information on nonlcon.	fgoalattain, fmincon, fminimax
ntheta	The number of semi-infinite constraints.	fseminf

Argument	Description	Used by Functions
options	A structure that defines options used by the optimization functions. For information about the options, see “Optimization Options Reference” on page 9-7 or the individual function reference pages.	All functions
semifcon	The function that computes the nonlinear inequality and equality constraints <i>and</i> the semi-infinite constraints. <code>semifcon</code> is the name of a function file or MEX-file. “Passing Extra Parameters” on page 2-53 explains how to parameterize <code>semifcon</code> , if necessary. See the function reference pages for <code>fseminf</code> for more information on <code>semifcon</code> .	<code>fseminf</code>
weight	A weighting vector to control the relative underattainment or overattainment of the objectives.	<code>fgoalattain</code>
xdata, ydata	The input data <code>xdata</code> and the observed output data <code>ydata</code> that are to be fitted to an equation.	<code>lsqcurvefit</code>
x0	Starting point (a scalar, vector or matrix). (For <code>fzero</code> , <code>x0</code> can also be a two-element vector representing a finite interval that is known to contain a zero.)	All functions except <code>fminbnd</code>
x1, x2	The interval over which the function is minimized.	<code>fminbnd</code>

Output Arguments

Argument	Description	Used by Functions
<code>attainfactor</code>	The attainment factor at the solution <code>x</code> .	<code>fgoalattain</code>
<code>exitflag</code>	<p>An integer identifying the reason the optimization algorithm terminated. See the function reference pages for descriptions of <code>exitflag</code> specific to each function, and “Exit Flags and Exit Messages” on page 3-3.</p> <p>You can also return a message stating why an optimization terminated by calling the optimization function with the output argument <code>output</code> and then displaying <code>output.message</code>.</p>	All functions
<code>fval</code>	The value of the objective function <code>fun</code> at the solution <code>x</code> .	<code>bintprog</code> , <code>fgoalattain</code> , <code>fminbnd</code> , <code>fmincon</code> , <code>fminimax</code> , <code>fminsearch</code> , <code>fminunc</code> , <code>fseminf</code> , <code>fsolve</code> , <code>fzero</code> , <code>linprog</code> , <code>quadprog</code>
<code>grad</code>	The value of the gradient of <code>fun</code> at the solution <code>x</code> . If <code>fun</code> does not compute the gradient, <code>grad</code> is a finite-differencing approximation of the gradient.	<code>fmincon</code> , <code>fminunc</code>
<code>hessian</code>	The value of the Hessian of <code>fun</code> at the solution <code>x</code> . For large-scale methods, if <code>fun</code> does not compute the Hessian, <code>hessian</code> is a finite-differencing approximation of the Hessian. For medium-scale methods, <code>hessian</code> is the value of the Quasi-Newton approximation to the Hessian at the solution <code>x</code> . See “Hessian” on page 3-28.	<code>fmincon</code> , <code>fminunc</code>

Argument	Description	Used by Functions
jacobian	The value of the Jacobian of fun at the solution x. If fun does not compute the Jacobian, jacobian is a finite-differencing approximation of the Jacobian.	lsqcurvefit, lsqnonlin, fsolve
lambda	The Lagrange multipliers at the solution x, see “Lagrange Multiplier Structures” on page 3-27. lambda is a structure where each field is for a different constraint type. For structure field names, see individual function descriptions. (For lsqnonneg, lambda is simply a vector, as lsqnonneg only handles one kind of constraint.)	fgoalattain, fmincon, fminimax, fseminf, linprog, lsqcurvefit, lsqlin, lsqnonlin, lsqnonneg, quadprog
maxfval	$\max\{\text{fun}(x)\}$ at the solution x.	fminimax
output	An output structure that contains information about the results of the optimization, see “Output Structures” on page 3-26. For structure field names, see individual function descriptions.	All functions
residual	The value of the residual at the solution x.	lsqcurvefit, lsqlin, lsqnonlin, lsqnonneg
resnorm	The value of the squared 2-norm of the residual at the solution x.	lsqcurvefit, lsqlin, lsqnonlin, lsqnonneg
x	The solution found by the optimization function. If exitflag > 0, then x is a solution; otherwise, x is the value of the optimization routine when it terminated prematurely.	All functions

Optimization Options Reference

In this section...
“Optimization Options” on page 9-7
“Output Function” on page 9-21
“Plot Functions” on page 9-30

Optimization Options

The following table describes optimization options. Create options using the `optimoptions` function, or `optimset` for `fminbnd`, `fminsearch`, `fzero`, or `lsqnonneg`.

See the individual function reference pages for information about available option values and defaults.

The default values for the options vary depending on which optimization function you call with `options` as an input argument. You can determine the default option values for any of the optimization functions by entering `optimoptions(@solvername)` or the equivalent `optimoptions('solvername')`. For example,

```
optimoptions(@fmincon)
```

returns a list of the options and the default values for the default trust-region-reflective `fmincon` algorithm. To find the default values for another `fmincon` algorithm, set the `Algorithm` option. For example,

```
opts = optimoptions(@fmincon,'Algorithm','sqp')
```

Optimization Options

Option Name	Description	Used by Functions
Algorithm	Chooses the algorithm used by the solver.	fmincon, fminunc, fsolve, linprog, lsqcurvefit, lsqin, lsqnonlin, quadprog
AlwaysHonorConstraints	The default 'bounds' ensures that bound constraints are satisfied at every iteration. Turn off by setting to 'none'.	fmincon
BranchingRule	Rule for choosing the component for branching: <ul style="list-style-type: none"> 'maxpscost' — The fractional component with maximum pseudocost. See “Branch and Bound” on page 6-161. 'mostfractional' — The component whose fractional part is closest to 1/2. 'maxfun' — The fractional component with maximal corresponding component in the absolute value of objective vector f. 	intlinprog
BranchStrategy	Strategy bintprog uses to select branch variable.	bintprog
CutGeneration	Level of cut generation (see “Cut Generation” on page 6-160): <ul style="list-style-type: none"> 'none' — No cuts. Makes CutGenerationMaxIter irrelevant. 'basic' — Normal cut generation. 'intermediate' — Use more cut types. 'advanced' — Use most cut types. 	intlinprog

Optimization Options (Continued)

Option Name	Description	Used by Functions
CutGenMaxIter	Number of passes through all cut generation methods before entering the branch-and-bound phase, an integer from 1 through 50. Disable cut generation by setting the CutGeneration option to 'none'.	intlinprog
DerivativeCheck	Compare user-supplied analytic derivatives (gradients or Jacobian, depending on the selected solver) to finite differencing derivatives.	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin
Diagnostics	Display diagnostic information about the function to be minimized or solved.	All but fminbnd, fminsearch, fzero, and lsqnonneg
DiffMaxChange	Maximum change in variables for finite differencing.	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin
DiffMinChange	Minimum change in variables for finite differencing.	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin

Optimization Options (Continued)

Option Name	Description	Used by Functions
Display	<p>Level of display.</p> <ul style="list-style-type: none">• 'off' displays no output.• 'iter' displays output at each iteration, and gives the default exit message.• 'iter-detailed' displays output at each iteration, and gives the technical exit message.• 'notify' displays output only if the function does not converge, and gives the default exit message.• 'notify-detailed' displays output only if the function does not converge, and gives the technical exit message.• 'final' displays just the final output, and gives the default exit message.• 'final-detailed' displays just the final output, and gives the technical exit message.	All. See the individual function reference pages for the values that apply.

Optimization Options (Continued)

Option Name	Description	Used by Functions
FinDiffRelStep	<p>Scalar or vector step size factor. When you set FinDiffRelStep to a vector v, forward finite differences δ are</p> $\delta = v \cdot \text{sign}(x) \cdot \max(\text{abs}(x), \text{TypicalX});$ <p>and central finite differences are</p> $\delta = v \cdot \max(\text{abs}(x), \text{TypicalX});$ <p>Scalar FinDiffRelStep expands to a vector. The default is <code>sqrt(eps)</code> for forward finite differences, and <code>eps^(1/3)</code> for central finite differences.</p>	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin
FinDiffType	<p>Finite differences, used to estimate gradients, are either 'forward' (the default), or 'central' (centered), which takes twice as many function evaluations but should be more accurate. 'central' differences might violate bounds during their evaluation in fmincon interior-point evaluations if the AlwaysHonorConstraints option is set to 'none'.</p>	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin

Optimization Options (Continued)

Option Name	Description	Used by Functions
FunValCheck	<p>Check whether objective function and constraints values are valid. 'on' displays an error when the objective function or constraints return a value that is complex, NaN, or Inf.</p> <hr/> <p>Note FunValCheck does not return an error for Inf when used with fminbnd, fminsearch, or fzero, which handle Inf appropriately.</p> <hr/> <p>'off' displays no error.</p>	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fsemif, fsolve, fzero, lsqcurvefit, lsqnonlin
GoalsExactAchieve	Specify the number of objectives required for the objective fun to equal the goal goal. Objectives should be partitioned into the first few elements of F.	fgoalattain
GradConstr	User-defined gradients for the nonlinear constraints.	fgoalattain, fmincon, fminimax
GradObj	User-defined gradients for the objective functions.	fgoalattain, fmincon, fminimax, fminunc, fsemif
HessFcn	Function handle to a user-supplied Hessian (see “Hessian” on page 10-49).	fmincon
Hessian	If 'user-supplied', function uses user-defined Hessian or Hessian information (when using HessMult), for the objective function. If 'off', function approximates the Hessian using finite differences.	fmincon, fminunc

Optimization Options (Continued)

Option Name	Description	Used by Functions
HessMult	Handle to a user-supplied Hessian multiply function. For <code>fmincon</code> , ignored unless Hessian is 'user-supplied' or 'on'.	<code>fmincon</code> , <code>fminunc</code> , <code>quadprog</code>
HessPattern	Sparsity pattern of the Hessian for finite differencing. The size of the matrix is n-by-n, where n is the number of elements in <code>x0</code> , the starting point.	<code>fmincon</code> , <code>fminunc</code>
HessUpdate	Quasi-Newton updating scheme.	<code>fminunc</code>
Heuristics	Algorithm for searching for feasible points (see “Heuristics for Finding Feasible Solutions” on page 6-161): <ul style="list-style-type: none"> • 'none' • 'rss' • 'round' • 'rins' 	<code>intlinprog</code>
HeuristicsMaxNodes	Strictly positive integer that bounds the number of nodes <code>intlinprog</code> can explore in its branch-and-bound search for feasible points. See “Heuristics for Finding Feasible Solutions” on page 6-161.	<code>intlinprog</code>
InitBarrierParam	Initial barrier value.	<code>fmincon</code>
InitialHessMatrix	Initial quasi-Newton matrix.	<code>fminunc</code>
InitialHessType	Initial quasi-Newton matrix type.	<code>fminunc</code>
InitTrustRegionRadius	Initial radius of the trust region.	<code>fmincon</code>

Optimization Options (Continued)

Option Name	Description	Used by Functions
IPPreprocess	<p>Types of integer preprocessing (see “Mixed-Integer Program Preprocessing” on page 6-159):</p> <ul style="list-style-type: none"> • 'none' — Use very few integer preprocessing steps. • 'basic' — Use a moderate number of integer preprocessing steps. • 'advanced' — Use all available integer preprocessing steps. 	intlinprog
Jacobian	<p>If 'on', function uses user-defined Jacobian or Jacobian information (when using JacobMult), for the objective function. If 'off', function approximates the Jacobian using finite differences.</p>	fsolve, lsqcurvefit, lsqnonlin
JacobMult	<p>User-defined Jacobian multiply function. Ignored unless Jacobian is 'on' for fsolve, lsqcurvefit, and lsqnonlin.</p>	fsolve, lsqcurvefit, lsqlin, lsqnonlin
JacobPattern	<p>Sparsity pattern of the Jacobian for finite differencing. The size of the matrix is m-by-n, where m is the number of values in the first argument returned by the user-specified function fun, and n is the number of elements in x0, the starting point.</p>	fsolve, lsqcurvefit, lsqnonlin
LargeScale Use Algorithm instead	<p>Use large-scale algorithm if possible.</p>	fminunc, fsolve, linprog, lsqcurvefit, lsqlin, lsqnonlin

Optimization Options (Continued)

Option Name	Description	Used by Functions
LPMaXiter	Strictly positive integer, the maximum number of simplex algorithm iterations per node during the branch-and-bound process.	intlinprog
LPPreprocess	Type of preprocessing for the solution to the relaxed linear program (see “Linear Program Preprocessing” on page 6-158): <ul style="list-style-type: none"> • 'none' — No preprocessing. • 'basic' — Use preprocessing. 	intlinprog
MaxFunEvals	Maximum number of function evaluations allowed.	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fsemint, fsolve, lsqcurvefit, lsqnonlin
MaxIter	Maximum number of iterations allowed.	All but fzero and lsqnonneg
MaxNodes	Strictly positive integer that is the maximum number of nodes the solver explores in its branch-and-bound process.	bintprog, intlinprog
MaxNumFeasPoints	Strictly positive integer. intlinprog stops if it finds MaxNumFeasPoints integer feasible points.	intlinprog
MaxPCGiter	Maximum number of iterations of preconditioned conjugate gradients method allowed.	fmincon, fminunc, fsolve, lsqcurvefit, lsqlin, lsqnonlin, quadprog

Optimization Options (Continued)

Option Name	Description	Used by Functions
MaxProjCGIter	A tolerance for the number of projected conjugate gradient iterations; this is an inner iteration, not the number of iterations of the algorithm.	fmincon
MaxRLPIter	Maximum number of iterations of linear programming relaxation method allowed.	bintprog
MaxSQPIter	Maximum number of iterations of sequential quadratic programming method allowed.	fgoalattain, fmincon, fminimax
MaxTime	Maximum amount of time in seconds allowed for the algorithm.	bintprog, intlinprog
MeritFunction	Use goal attainment/minimax merit function (multiobjective) vs. fmincon (single objective).	fgoalattain, fminimax
MinAbsMax	Number of $F(x)$ to minimize the worst case absolute values.	fminimax
NodeDisplayInterval	Node display interval for bintprog.	bintprog
NodeSearchStrategy	Search strategy that bintprog uses.	bintprog
NodeSelection	Choose the node to explore next. <ul style="list-style-type: none"> 'simplebestproj' — Best projection. See “Branch and Bound” on page 6-161. 'minobj' — Explore the node with the minimum objective function. 'mininfeas' — Explore the node with the minimal sum of integer infeasibilities. See “Branch and Bound” on page 6-161. 	intlinprog
ObjectiveCutoff	Real greater than $-\text{Inf}$. The default is Inf .	intlinprog

Optimization Options (Continued)

Option Name	Description	Used by Functions
ObjectiveLimit	If the objective function value goes below ObjectiveLimit and the iterate is feasible, then the iterations halt.	fmincon, fminunc, quadprog
OutputFcn	Specify one or more user-defined functions that the optimization function calls at each iteration. See “Output Function” on page 9-21.	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fsemif, fsolve, fzero, lsqcurvefit, lsqnonlin
PlotFcns	<p>Plots various measures of progress while the algorithm executes, select from predefined plots or write your own.</p> <ul style="list-style-type: none"> • @optimplotx plots the current point • @optimplotfuncount plots the function count • @optimplotfval plots the function value • @optimplotconstrviolation plots the maximum constraint violation • @optimplotresnorm plots the norm of the residuals • @optimplotfirstorderopt plots the first-order of optimality • @optimplotstepsize plots the step size <p>See “Plot Functions” on page 9-30.</p>	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fsemif, fsolve, fzero, lsqcurvefit, lsqnonlin. See the individual function reference pages for the values that apply.

Optimization Options (Continued)

Option Name	Description	Used by Functions
PrecondBandWidth	Upper bandwidth of preconditioner for PCG. Setting to 'Inf' uses a direct factorization instead of CG.	fmincon, fminunc, fsolve, lsqcurvefit, lsqlin, lsqnonlin, quadprog
RelLineSrchBnd	Relative bound on line search step length.	fgoalattain, fmincon, fminimax, fseminf
RelLineSrchBndDuration	Number of iterations for which the bound specified in RelLineSrchBnd should be active.	fgoalattain, fmincon, fminimax, fseminf
RelObjThreshold	Nonnegative real. intlinprog changes the current feasible solution only when it locates another with an objective function value that is at least RelObjThreshold lower: $(fold - fnew)/(1 + fold) > RelObjThreshold$.	intlinprog
RootLPAlgorithm	Algorithm for solving linear programs: <ul style="list-style-type: none"> 'dual-simplex' — Dual simplex algorithm 'primal-simplex' — Primal simplex algorithm 	intlinprog
RootLPMaxIter	Nonnegative integer that is the maximum number of simplex algorithm iterations to solve the initial linear programming problem.	intlinprog

Optimization Options (Continued)

Option Name	Description	Used by Functions
ScaleProblem	<p>For fmincon interior-point and sqp algorithms, 'obj-and-constr' causes the algorithm to normalize all constraints and the objective function by their initial values. Disable by setting to the default 'none'.</p> <p>For the other solvers, when using the Algorithm option 'levenberg-marquardt', setting the ScaleProblem option to 'jacobian' sometimes helps the solver on badly-scaled problems.</p>	fmincon, fsolve, lsqcurvefit, lsqnonlin, quadprog
Simplex Use Algorithm instead	If 'on', function uses the simplex algorithm.	linprog
SubproblemAlgorithm	Determines how the iteration step is calculated.	fmincon
TolCon	Tolerance on the constraint violation.	bintprog, fgoalattain, fmincon, fminimax, fseminf, intlinprog, quadprog
TolConSQP	Constraint violation tolerance for the inner SQP iteration.	fgoalattain, fmincon, fminimax, fseminf
TolFun	Termination tolerance on the function value.	bintprog, fgoalattain, fmincon, fminimax, fminsearch, fminunc, fseminf, fsolve, linprog (interior-point only), lsqcurvefit, lsqlin (trust-region-reflective only), lsqnonlin, quadprog

Optimization Options (Continued)

Option Name	Description	Used by Functions
TolFunLP	Nonnegative real where reduced costs must exceed TolFunLP for a variable to be taken into the basis.	intlinprog
TolGapAbs	Nonnegative real. intlinprog stops if the difference between the internally calculated upper (U) and lower (L) bounds on the objective function is less than or equal to TolGapAbs: $U - L \leq \text{TolGapAbs}$.	intlinprog
TolGapRel	Real from 0 through 1. intlinprog stops if the relative difference between the internally calculated upper (U) and lower (L) bounds on the objective function is less than or equal to TolGapRel: $(U - L) / (\text{abs}(U) + 1) \leq \text{TolGapRel}$.	intlinprog
TolInteger	Real from 1e-6 through 1e-3, where the maximum deviation from integer that a component of the solution x can have and still be considered an integer. TolInteger is not a stopping criterion.	intlinprog
TolPCG	Termination tolerance on the PCG iteration.	fmincon, fminunc, fsolve, lsqcurvefit, lsqlin, lsqnonlin, quadprog
TolProjCG	A relative tolerance for projected conjugate gradient algorithm; this is for an inner iteration, not the algorithm iteration.	fmincon
TolProjCGAbs	Absolute tolerance for projected conjugate gradient algorithm; this is for an inner iteration, not the algorithm iteration.	fmincon

Optimization Options (Continued)

Option Name	Description	Used by Functions
To1RLPFun	Termination tolerance on the function value of a linear programming relaxation problem.	bintprog
TolX	Termination tolerance on x .	All functions except the medium-scale algorithms for linprog, lsqlin, and quadprog
TolXInteger	Tolerance within which bintprog considers the value of a variable to be an integer.	bintprog
TypicalX	Array that specifies typical magnitude of array of parameters x . The size of the array is equal to the size of x_0 , the starting point. Primarily for scaling finite differences for gradient estimation.	fgoalattain, fmincon, fminimax, fminunc, fsolve, lsqcurvefit, lsqlin, lsqnonlin, quadprog
UseParallel	When true, applicable solvers estimate gradients in parallel. Disable by setting to false.	fgoalattain, fmincon, fminimax.

Output Function

The `Outputfcn` field of options specifies one or more functions that an optimization function calls at each iteration. Typically, you might use an output function to plot points at each iteration or to display optimization quantities from the algorithm. Using an output function you can view, but not set, optimization quantities. To set up an output function, do the following:

- 1** Write the output function as a function file or local function.
- 2** Use `optimoptions` to set the value of `Outputfcn` to be a function handle, that is, the name of the function preceded by the `@` sign. For example, if the output function is `outfun.m`, the command

```
options = optimoptions(@solvername, 'OutputFcn', @outfun);
```

specifies `OutputFcn` to be the handle to `outfun`. To specify more than one output function, use the syntax

```
options = optimoptions(@solvername, 'OutputFcn', {@outfun, @outfun2});
```

3 Call the optimization function with `options` as an input argument.

See “Output Functions” on page 3-37 for an example of an output function.

“Passing Extra Parameters” on page 2-53 explains how to parameterize the output function `OutputFcn`, if necessary.

Structure of the Output Function

The function definition line of the output function has the following form:

```
stop = outfun(x, optimValues, state)
```

where

- `x` is the point computed by the algorithm at the current iteration.
- `optimValues` is a structure containing data from the current iteration. “Fields in `optimValues`” on page 9-22 describes the structure in detail.
- `state` is the current state of the algorithm. “States of the Algorithm” on page 9-29 lists the possible values.
- `stop` is a flag that is `true` or `false` depending on whether the optimization routine should quit or continue. See “Stop Flag” on page 9-29 for more information.

The optimization function passes the values of the input arguments to `outfun` at each iteration.

Fields in `optimValues`

The following table lists the fields of the `optimValues` structure. A particular optimization function returns values for only some of these fields. For each field, the Returned by Functions column of the table lists the functions that return the field.

Corresponding Output Arguments. Some of the fields of `optimValues` correspond to output arguments of the optimization function. After the final iteration of the optimization algorithm, the value of such a field equals the corresponding output argument. For example, `optimValues.fval` corresponds to the output argument `fval`. So, if you call `fmincon` with an output function and return `fval`, the final value of `optimValues.fval` equals `fval`. The Description column of the following table indicates the fields that have a corresponding output argument.

Command-Line Display. The values of some fields of `optimValues` are displayed at the command line when you call the optimization function with the `Display` field of `options` set to `'iter'`, as described in “Iterative Display” on page 3-17. For example, `optimValues.fval` is displayed in the `f(x)` column. The Command-Line Display column of the following table indicates the fields that you can display at the command line.

Some `optimValues` fields apply only to specific algorithms:

- AS — active-set
- D — trust-region-dogleg
- IP — interior-point
- LM — levenberg-marquardt
- Q — quasi-newton
- SQP — sqp
- TR — trust-region
- TRR — trust-region-reflective

optimValues Fields

OptimValues Field (optimValues.field)	Description	Returned by Functions	Command-Line Display
attainfactor	Attainment factor for multiobjective problem. For details, see “Goal Attainment Method” on page 6-237.	fgoalattain	None
cgiterations	Number of conjugate gradient iterations at current optimization iteration.	fmincon (IP, TRR), fsolve (TRR), lsqcurvefit (TRR), lsqnonlin (TRR)	CG-iterations See “Iterative Display” on page 3-17.
constrviolation	Maximum constraint violation.	fgoalattain, fmincon, fminimax, fseminf	Max constraint or Feasibility See “Iterative Display” on page 3-17.
degenerate	Measure of degeneracy. A point is <i>degenerate</i> if the partial derivative with respect to one of the variables is 0 at the point. A bound constraint is active for that variable at the point. See “Degeneracy” on page 9-27.	fmincon (TRR), lsqcurvefit (TRR), lsqnonlin (TRR)	None
directionalderivative	Directional derivative in the search direction.	fgoalattain, fmincon (AS), fminimax, fminunc (Q), fseminf, fsolve (LM), lsqcurvefit (LM), lsqnonlin (LM)	Directional derivative See “Iterative Display” on page 3-17.

optimValues Fields (Continued)

OptimValues Field (optimValues.field)	Description	Returned by Functions	Command-Line Display
firstorderopt	First-order optimality (depends on algorithm). Final value equals optimization function output <code>output.firstorderopt</code> .	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin	First-order optimality See “Iterative Display” on page 3-17.
funccount	Cumulative number of function evaluations. Final value equals optimization function output <code>output.funcCount</code> .	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fsolve, fzero, fseminf, lsqcurvefit, lsqnonlin	F-count or Func-count See “Iterative Display” on page 3-17.
fval	Function value at current point. Final value equals optimization function output <code>fval</code> .	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fseminf, fsolve, fzero	f(x) See “Iterative Display” on page 3-17.
gradient	Current gradient of objective function — either analytic gradient if you provide it or finite-differencing approximation. Final value equals optimization function output <code>grad</code> .	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin	None
iteration	Iteration number — starts at 0. Final value equals optimization function output <code>output.iterations</code> .	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fsolve, fseminf, fzero, lsqcurvefit, lsqnonlin	Iteration See “Iterative Display” on page 3-17.

optimValues Fields (Continued)

OptimValues Field (optimValues.field)	Description	Returned by Functions	Command-Line Display
lambda	The Levenberg-Marquardt parameter, <code>lambda</code> , at the current iteration. See “Levenberg-Marquardt Method” on page 6-205.	<code>fsolve</code> (LM), <code>lsqcurvefit</code> (LM), <code>lsqnonlin</code> (LM)	Lambda
maxfval	Maximum function value	<code>fminimax</code>	None
positivedefinite	0 if algorithm detects negative curvature while computing Newton step. 1 otherwise.	<code>fmincon</code> (TRR), <code>fminunc</code> (TRR), <code>fsolve</code> (TRR), <code>lsqcurvefit</code> (TRR), <code>lsqnonlin</code> (TRR)	None
procedure	Procedure messages.	<code>fgoalattain</code> , <code>fminbnd</code> , <code>fmincon</code> (AS), <code>fminimax</code> , <code>fminsearch</code> , <code>fseminf</code> , <code>fzero</code>	Procedure See “Iterative Display” on page 3-17.
ratio	Ratio of change in the objective function to change in the quadratic approximation.	<code>fmincon</code> (TRR), <code>fsolve</code> (TRR), <code>lsqcurvefit</code> (TRR), <code>lsqnonlin</code> (TRR)	None
residual	The residual vector. For <code>fsolve</code> , <code>residual</code> means the 2-norm of the residual squared.	<code>lsqcurvefit</code> , <code>lsqnonlin</code> , <code>fsolve</code>	Residual See “Iterative Display” on page 3-17.
resnorm	2-norm of the residual squared.	<code>lsqcurvefit</code> , <code>lsqnonlin</code>	Resnorm See “Iterative Display” on page 3-17.

optimValues Fields (Continued)

OptimValues Field (optimValues.field)	Description	Returned by Functions	Command-Line Display
searchdirection	Search direction.	fgoalattain, fmincon (AS, SQP), fminimax, fminunc (Q), fseminf, fsolve (LM), lsqcurvefit (LM), lsqnonlin (LM)	None
stepaccept	Status of the current trust-region step. Returns true if the current trust-region step was successful, and false if the trust-region step was unsuccessful.	fsolve (D)	None
stepsize	Current step size (displacement in x). Final value equals optimization function output <code>output.stepsize</code> .	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin	Step-size or Norm of Step See “Iterative Display” on page 3-17.
trustregionradius	Radius of trust region.	fmincon (IP, TRR), fminunc (TR), fsolve (D, TRR), lsqcurvefit (TRR), lsqnonlin (TRR)	Trust-region radius See “Iterative Display” on page 3-17.

Degeneracy. The value of the field `degenerate`, which measures the degeneracy of the current optimization point x , is defined as follows. First, define a vector r , of the same size as x , for which $r(i)$ is the minimum distance from $x(i)$ to the i th entries of the lower and upper bounds, `lb` and `ub`. That is,

$$r = \min(\text{abs}(\text{ub}-x, x-\text{lb}))$$

Then the value of `degenerate` is the minimum entry of the vector $r + \text{abs}(\text{grad})$, where `grad` is the gradient of the objective function. The value of `degenerate` is 0 if there is an index `i` for which both of the following are true:

- `grad(i) = 0`
- `x(i)` equals the *i*th entry of either the lower or upper bound.

States of the Algorithm

The following table lists the possible values for `state`:

State	Description
'init'	The algorithm is in the initial state before the first iteration.
'interrupt'	The algorithm is in some computationally expensive part of the iteration. In this state, the output function can interrupt the current iteration of the optimization. At this time, the values of <code>x</code> and <code>optimValues</code> are the same as at the last call to the output function in which <code>state=='iter'</code> .
'iter'	The algorithm is at the end of an iteration.
'done'	The algorithm is in the final state after the last iteration.

The following code illustrates how the output function might use the value of `state` to decide which tasks to perform at the current iteration:

```
switch state
    case 'iter'
        % Make updates to plot or guis as needed
    case 'interrupt'
        % Probably no action here. Check conditions to see
        % whether optimization should quit.
    case 'init'
        % Setup for plots or guis
    case 'done'
        % Cleanup of plots, guis, or final plot
otherwise
end
```

Stop Flag

The output argument `stop` is a flag that is `true` or `false`. The flag tells the optimization function whether the optimization should quit or continue. The following examples show typical ways to use the stop flag.

Stopping an Optimization Based on Data in `optimValues`. The output function can stop an optimization at any iteration based on the current data in `optimValues`. For example, the following code sets `stop` to `true` if the directional derivative is less than `.01`:

```
function stop = outfun(x,optimValues,state)
stop = false;
% Check if directional derivative is less than .01.
if optimValues.directionalderivative < .01
    stop = true;
end
```

Stopping an Optimization Based on GUI Input. If you design a GUI to perform optimizations, you can make the output function stop an optimization when a user clicks a **Stop** button on the GUI. The following code shows how to do this, assuming that the **Stop** button callback stores the value `true` in the `optimstop` field of a `handles` structure called `hObject`:

```
function stop = outfun(x,optimValues,state)
stop = false;
% Check if user has requested to stop the optimization.
stop = getappdata(hObject,'optimstop');
```

Plot Functions

The `PlotFcns` field of the `options` structure specifies one or more functions that an optimization function calls at each iteration to plot various measures of progress while the algorithm executes. The structure of a plot function is the same as that for an output function. For more information on writing and calling a plot function, see “Output Function” on page 9-21. For an example of using built-in plot functions, “Using a Plot Function” on page 3-31.

To view a predefined plot function listed for `PlotFcns`, you can open it in the MATLAB Editor. For example, to view the file corresponding to the norm of residuals, enter:

```
edit optimplotresnorm.m
```

You can use any predefined plot function as a template for writing a custom plot function.

Functions — Alphabetical List

Purpose Solve binary integer programming problems

Equation Solves binary integer programming problems of the form

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ x \text{ binary.} \end{cases}$$

f , b , and beq are vectors, A and Aeq are matrices, and the solution x is required to be a binary integer vector—that is, its entries can only take on the values 0 or 1.

Syntax

```
x = bintprog(f)
x = bintprog(f,A,b)
x = bintprog(f,A,b,Aeq,beq)
x = bintprog(f,A,b,Aeq,beq,x0)
x = bintprog(f,A,b,Aeq,Beq,x0,options)
x = bintprog(problem)
[x,fval] = bintprog(...)
[x,fval,exitflag] = bintprog(...)
[x,fval,exitflag,output] = bintprog(...)
```

Description $x = \text{bintprog}(f)$ solves the binary integer programming problem

$$\min_x f^T x.$$

$x = \text{bintprog}(f,A,b)$ solves the binary integer programming problem

$$\min_x f^T x \text{ such that } A \cdot x \leq b.$$

$x = \text{bintprog}(f,A,b,Aeq,beq)$ solves the preceding problem with the additional equality constraint.

$$Aeq \cdot x = beq.$$

`x = bintprog(f,A,b,Aeq,beq,x0)` sets the starting point for the algorithm to `x0`. If `x0` is not in the feasible region, `bintprog` uses the default initial point.

`x = bintprog(f,A,b,Aeq,Beq,x0,options)` minimizes with the default optimization options replaced by values in `options`, which you can create using the function `optimoptions`.

`x = bintprog(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-3.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = bintprog(...)` returns `fval`, the value of the objective function at `x`.

`[x,fval,exitflag] = bintprog(...)` returns `exitflag` that describes the exit condition of `bintprog`. See “Output Arguments” on page 10-4.

`[x,fval,exitflag,output] = bintprog(...)` returns a structure `output` that contains information about the optimization. See “Output Arguments” on page 10-4.

Input Arguments

The following table lists the input arguments for `bintprog`. “Function Arguments” on page 9-2 contains general descriptions of input arguments for optimization functions.

<code>f</code>	Vector containing the coefficients of the linear objective function.
<code>A</code>	Matrix containing the coefficients of the linear inequality constraints $Ax \leq b$.
<code>b</code>	Vector corresponding to the right-hand side of the linear inequality constraints.
<code>Aeq</code>	Matrix containing the coefficients of the linear equality constraints $Aeqx = beq$.

<code>beq</code>		Vector containing the constants of the linear equality constraints.
<code>x0</code>		Initial point for the algorithm.
<code>options</code>		Options for the algorithm.
<code>problem</code>	<code>f</code>	Linear objective function vector <code>f</code>
	<code>Aineq</code>	Matrix for linear inequality constraints
	<code>bineq</code>	Vector for linear inequality constraints
	<code>Aeq</code>	Matrix for linear equality constraints
	<code>beq</code>	Vector for linear equality constraints
	<code>x0</code>	Initial point for <code>x</code>
	<code>solver</code>	'bintprog'
	<code>options</code>	Options created with <code>optimoptions</code>

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `bintprog`. This section provides specific details for the arguments `exitflag`, `output`, and iterative display.

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.
1	Function converged to a solution <code>x</code> .
0	Number of iterations exceeded <code>options.MaxIter</code> .
-2	The problem is infeasible.
-4	Number of searched nodes exceeded <code>options.MaxNodes</code> .
-5	Search time exceeded <code>options.MaxTime</code> .

	-6	Number of iterations the LP-solver performed at a node to solve the LP-relaxation problem exceeded <code>options.MaxRLP</code> .
output	Structure containing information about the optimization. The fields of the structure are	
	<code>iterations</code>	Number of iterations taken
	<code>nodes</code>	Number of nodes searched
	<code>time</code>	Execution time of the algorithm
	<code>algorithm</code>	Optimization algorithm used
	<code>branchStrategy</code>	Strategy used to select branch variable—see “Options” on page 10-6
	<code>nodeSearchStrategy</code>	Strategy used to select next node in search tree—see “Options” on page 10-6
	<code>message</code>	Exit message
Iterative display	bintprog-specific iterative display:	
	<code>Best lower bound on obj</code>	Objective function value of LP relaxation problem that gives the best current lower bound on the final objective function value.
	<code>Explored nodes</code>	Cumulative number of explored nodes.
	<code>Obj of best integer point</code>	Objective function value of the best integer point found so far. This is an upper bound for the final objective function value.

Obj of LP relaxation	Objective function value of the linear programming (LP) relaxation problem.
Relative gap between bounds	$\frac{100(b - a)}{ b + 1},$ <p>where</p> <ul style="list-style-type: none">• b is the objective function value of the best integer point.• a is the best lower bound on the objective function value.
Unexplored nodes	Number of nodes that have been set up but not yet explored.

Options

Optimization options used by bintprog. Use `optoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information.

BranchStrategy	Strategy the algorithm uses to select the branch variable in the search tree—see “Branching” on page 6-148. The choices are <ul style="list-style-type: none">• 'mininfeas' — Choose the variable with the minimum integer infeasibility (the variable whose value is closest to 0 or 1, but not equal to 0 or 1).• 'maxinfeas' — Choose the variable with the maximum integer infeasibility (the
----------------	--

	variable whose value is closest to 0.5 (default)).
Diagnostics	Display diagnostic information about the function. The choices are 'on' or the default, 'off'.
Display	Level of display: <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'iter' displays output at each iteration, and gives the default exit message.• 'iter-detailed' displays output at each iteration, and gives the technical exit message.• 'final' (default) displays just the final output, and gives the default exit message.• 'final-detailed' displays just the final output, and gives the technical exit message.
MaxIter	Maximum number of iterations allowed (a positive integer). The default is $100000 * \text{numberOfVariables}$
MaxNodes	Maximum number of solutions, or nodes, the function searches (a positive integer). The default is $1000 * \text{numberOfVariables}$
MaxRLPIter	Maximum number of iterations the LP-solver performs to solve the LP-relaxation problem at each node (a positive integer). The default is $100 * \text{numberOfVariables}$

MaxTime	Maximum amount of CPU time in seconds the function runs (a positive scalar). The default is 7200.
NodeDisplayInterval	Node display interval (a positive integer). Gives the number of nodes to search between reporting to the iterative display. The default is 20.
NodeSearchStrategy	Strategy the algorithm uses to select the next node to search in the search tree—see “Branching” on page 6-148. The choices are: <ul style="list-style-type: none">• 'df' — Depth-first search strategy. At each node in the search tree, if there is a child node one level down in the tree that has not already been explored, the algorithm chooses one such child to search. Otherwise, the algorithm moves to the node one level up in the tree and chooses a child node one level down from that node.• 'bn' — Best-node search strategy, which chooses the node with lowest bound on the objective function (the default).
TolFun	Termination tolerance on the function value (a positive scalar). The default is $1.0e-3$.
TolXInteger	Tolerance within which the value of a variable is considered to be integral (a positive scalar). The default is $1.0e-8$.
TolRLPFun	Termination tolerance on the function value of a linear programming relaxation problem (a positive scalar). The default is $1.0e-6$.

Examples

To minimize the function

$$f(x) = -9x_1 - 5x_2 - 6x_3 - 4x_4,$$

subject to the constraints

$$\begin{bmatrix} 6 & 3 & 5 & 2 \\ 0 & 0 & 1 & 1 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} \leq \begin{bmatrix} 9 \\ 1 \\ 0 \\ 0 \end{bmatrix},$$

where x_1 , x_2 , x_3 , and x_4 are binary integers, enter the following commands:

```
f = [-9; -5; -6; -4];
A = [6 3 5 2; 0 0 1 1; -1 0 1 0; 0 -1 0 1];
b = [9; 1; 0; 0];
x = bintprog(f,A,b)
```

Optimization terminated.

```
x =
     1
     1
     0
     0
```

Tips

- `intlinprog` solves more problems than `bintprog`, and has better performance. To update your existing `bintprog` code to use `intlinprog`, make the following changes:
 - Set `intcon` to `1:numVars`, where `numVars` is the number of variables in your problem.
 - Set `lb` to `zeros(numVars,1)`.
 - Set `ub` to `ones(numVars,1)`.

- Update any relevant options. Use `optimoptions` to create options for `intlinprog`.
- Change your call to `bintprog` as follows:

```
[x,fval,exitflag,output] = bintprog(f,A,b,Aeq,Beq,x0,options)
```

```
% Change your call to:
```

```
[x,fval,exitflag,output] = intlinprog(f,intcon,A,b,Aeq,Beq,lb,ub,opt
```

Algorithms

`bintprog` uses a linear programming (LP)-based branch-and-bound algorithm to solve binary integer programming problems. The algorithm searches for an optimal solution to the binary integer programming problem by solving a series of *LP-relaxation* problems, in which the binary integer requirement on the variables is replaced by the weaker constraint $0 \leq x \leq 1$. The algorithm

- Searches for a binary integer feasible solution
- Updates the best binary integer feasible point found so far as the search tree grows
- Verifies that no better integer feasible solution is possible by solving a series of linear programming problems

For more information, see “`bintprog` Algorithm” on page 6-148

References

[1] Wolsey, Laurence A., *Integer Programming*, John Wiley & Sons, 1998.

[2] Nemhauser, George L. and Laurence A. Wolsey, *Integer and Combinatorial Optimization*, John Wiley & Sons, 1988.

[3] Hillier, Frederick S. and Lieberman Gerald J., *Introduction to Operations Research*, McGraw-Hill, 2001.

See Also

`intlinprog` | `optimoptions` | `optimtool`

How To

- “Binary Integer Programming Algorithms” on page 6-148

- “Optimal Investments Via Binary Integer Programming” on page 6-151

color

Purpose Column partition for sparse finite differences

Syntax `group = color(J,P)`

Description `group = color(J,P)` returns a partition of the column corresponding to a coloring of the column-intersection graph. `GROUP(I) = J` means column `I` is colored `J`.

All columns belonging to a color can be estimated in a single finite difference.

Purpose Solve multiobjective goal attainment problems

Equation Finds the minimum of a problem specified by

$$\text{minimize}_{x,\gamma} \gamma \quad \text{such that} \quad \left\{ \begin{array}{l} F(x) - \text{weight} \cdot \gamma \leq \text{goal} \\ c(x) \leq 0 \\ \text{ceq}(x) = 0 \\ A \cdot x \leq b \\ \text{Aeq} \cdot x = \text{beq} \\ lb \leq x \leq ub. \end{array} \right.$$

weight, *goal*, *b*, and *beq* are vectors, *A* and *Aeq* are matrices, and *c(x)*, *ceq(x)*, and *F(x)* are functions that return vectors. *F(x)*, *c(x)*, and *ceq(x)* can be nonlinear functions.

x, *lb*, and *ub* can be passed as vectors or matrices; see “Matrix Arguments” on page 2-32.

Syntax

```
x = fgoalattain(fun,x0,goal,weight)
x = fgoalattain(fun,x0,goal,weight,A,b)
x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq)
x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq,lb,ub)
x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq,lb,ub,nonlcon)
x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq,lb,ub,nonlcon,
    ... options)
x = fgoalattain(problem)
[x,fval] = fgoalattain(...)
[x,fval,attainfactor] = fgoalattain(...)
[x,fval,attainfactor,exitflag] = fgoalattain(...)
[x,fval,attainfactor,exitflag,output] = fgoalattain(...)
[x,fval,attainfactor,exitflag,output,lambda] = fgoalattain(...)
```

Description

`fgoalattain` solves the goal attainment problem, which is one formulation for minimizing a multiobjective optimization problem.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective functions and nonlinear constraint functions, if necessary.

`x = fgoalattain(fun,x0,goal,weight)` tries to make the objective functions supplied by `fun` attain the goals specified by `goal` by varying `x`, starting at `x0`, with `weight` specified by `weight`.

`x = fgoalattain(fun,x0,goal,weight,A,b)` solves the goal attainment problem subject to the linear inequalities $A*x \leq b$.

`x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq)` solves the goal attainment problem subject to the linear equalities $Aeq*x = beq$ as well. Set `A = []` and `b = []` if no inequalities exist.

`x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq,lb,ub)` defines a set of lower and upper bounds on the design variables in `x`, so that the solution is always in the range $lb \leq x \leq ub$.

Note See “Iterations Can Violate Constraints” on page 2-34.

`x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq,lb,ub,nonlcon)` subjects the goal attainment problem to the nonlinear inequalities $c(x)$ or nonlinear equality constraints $ceq(x)$ defined in `nonlcon`. `fgoalattain` optimizes such that $c(x) \leq 0$ and $ceq(x) = 0$. Set `lb = []` and/or `ub = []` if no bounds exist.

`x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq,lb,ub,nonlcon, ... options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options.

`x = fgoalattain(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-15.

Create the problem structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = fgoalattain(...)` returns the values of the objective functions computed in `fun` at the solution `x`.

`[x,fval,attainfactor] = fgoalattain(...)` returns the attainment factor at the solution `x`.

`[x,fval,attainfactor,exitflag] = fgoalattain(...)` returns a value `exitflag` that describes the exit condition of `fgoalattain`.

`[x,fval,attainfactor,exitflag,output] = fgoalattain(...)` returns a structure `output` that contains information about the optimization.

`[x,fval,attainfactor,exitflag,output,lambda] = fgoalattain(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the output `fval` is `[]`.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `fgoalattain`. This section provides function-specific details for `fun`, `goal`, `nonlcon`, `options`, `weight`, and `problem`:

`fun` The function to be minimized. `fun` is a function that accepts a vector `x` and returns a vector `F`, the objective functions evaluated at `x`. The function `fun` can be specified as a function handle for a function file:

```
x = fgoalattain(@myfun,x0,goal,weight)
```

where `myfun` is a MATLAB function such as

```
function F = myfun(x)
F = ...           % Compute function values at x.
```

`fun` can also be a function handle for an anonymous function.

```
x = fgoalattain(@(x)sin(x.*x),x0,goal,weight);
```

If the user-defined values for `x` and `F` are matrices, they are converted to a vector using linear indexing.

To make an objective function as near as possible to a goal value, (i.e., neither greater than nor less than) use `optimoptions` to set the `GoalsExactAchieve` option to the number of objectives required to be in the neighborhood of the goal values. Such objectives *must* be partitioned into the first elements of the vector `F` returned by `fun`.

If the gradient of the objective function can also be computed *and* the `GradObj` option is `'on'`, as set by

```
options = optimoptions('fgoalattain','GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `G`, a matrix, at `x`. The gradient consists of the partial derivative dF/dx of each `F` at the point `x`. If `F` is a vector of length `m` and `x` has length `n`, where `n` is the length of `x0`, then the gradient `G`

of $F(x)$ is an n -by- m matrix where $G(i, j)$ is the partial derivative of $F(j)$ with respect to $x(i)$ (i.e., the j th column of G is the gradient of the j th objective function $F(j)$).

Note Setting `GradObj` to 'on' is effective only when there is no nonlinear constraint, or when the nonlinear constraint has `GradConstr` set to 'on' as well. This is because internally the objective is folded into the constraints, so the solver needs both gradients (objective and constraint) supplied in order to avoid estimating a gradient.

`goal` Vector of values that the objectives attempt to attain. The vector is the same length as the number of objectives `F` returned by `fun`. `fgoalattain` attempts to minimize the values in the vector `F` to attain the goal values given by `goal`.

`nonlcon` The function that computes the nonlinear inequality constraints $c(x) \leq 0$ and the nonlinear equality constraints $ceq(x) = 0$. The function `nonlcon` accepts a vector `x` and returns two vectors `c` and `ceq`. The vector `c` contains the nonlinear inequalities evaluated at `x`, and `ceq` contains the nonlinear equalities evaluated at `x`. The function `nonlcon` can be specified as a function handle.

```
x = fgoalattain(@myfun,x0,goal,weight,A,b,Aeq,beq,...
               lb,ub,@mycon)
```

where `mycon` is a MATLAB function such as

```
function [c,ceq] = mycon(x)
c = ...           % compute nonlinear inequalities at x.
ceq = ...         % compute nonlinear equalities at x.
```

If the gradients of the constraints can also be computed *and* the GradConstr option is 'on', as set by

```
options = optimoptions('fgoalattain','GradConstr','on')
```

then the function nonlcon must also return, in the third and fourth output arguments, GC, the gradient of $c(x)$, and GCEq, the gradient of $ceq(x)$. “Nonlinear Constraints” on page 2-37 explains how to “conditionalize” the gradients for use in solvers that do not accept supplied gradients.

If nonlcon returns a vector c of m components and x has length n , where n is the length of x_0 , then the gradient GC of $c(x)$ is an n -by- m matrix, where $GC(i, j)$ is the partial derivative of $c(j)$ with respect to $x(i)$ (i.e., the j th column of GC is the gradient of the j th inequality constraint $c(j)$). Likewise, if ceq has p components, the gradient GCEq of $ceq(x)$ is an n -by- p matrix, where $GCEq(i, j)$ is the partial derivative of $ceq(j)$ with respect to $x(i)$ (i.e., the j th column of GCEq is the gradient of the j th equality constraint $ceq(j)$).

Note Setting GradConstr to 'on' is effective only when GradObj is set to 'on' as well. This is because internally the objective is folded into the constraint, so the solver needs both gradients (objective and constraint) supplied in order to avoid estimating a gradient.

Note Because Optimization Toolbox functions only accept inputs of type `double`, user-supplied objective and nonlinear constraint functions must return outputs of type `double`.

“Passing Extra Parameters” on page 2-53 explains how to parameterize the nonlinear constraint function `nonlcon`, if necessary.

`options` “Options” on page 10-22 provides the function-specific details for the `options` values.

`weight` A weighting vector to control the relative underattainment or overattainment of the objectives in `fgoalattain`. When the values of `goal` are *all nonzero*, to ensure the same percentage of under- or overattainment of the active objectives, set the weighting function to `abs(goal)`. (The active objectives are the set of objectives that are barriers to further improvement of the goals at the solution.)

Note Setting a component of the `weight` vector to zero will cause the corresponding goal constraint to be treated as a hard constraint rather than as a goal constraint. An alternative method to set a hard constraint is to use the input argument `nonlcon`.

When the weighting function `weight` is positive, `fgoalattain` attempts to make the objectives less than the goal values. To make the objective functions greater than the goal values, set `weight` to be negative rather than positive. To make an objective function as near as possible to a goal value, use the `GoalsExactAchieve`

fgoalattain

option and put that objective as the first element of the vector returned by fun (see the preceding description of fun and options).

problem	objective	Vector of objective functions
	x0	Initial point for x
	goal	Goals to attain
	weight	Relative importance factors of goals
	Aineq	Matrix for linear inequality constraints
	bineq	Vector for linear inequality constraints
	Aeq	Matrix for linear equality constraints
	beq	Vector for linear equality constraints
	lb	Vector of lower bounds
	ub	Vector of upper bounds
	nonlcon	Nonlinear constraint function
	solver	'fgoalattain'
	options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by fgoalattain. This section provides function-specific details for attainfactor, exitflag, lambda, and output:

<code>attainfactor</code>	The amount of over- or underachievement of the goals. <code>attainfactor</code> contains the value of γ at the solution. If <code>attainfactor</code> is negative, the goals have been overachieved; if <code>attainfactor</code> is positive, the goals have been underachieved.
<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.
1	Function converged to a solutions x .
4	Magnitude of the search direction less than the specified tolerance and constraint violation less than <code>options.TolCon</code>
5	Magnitude of directional derivative less than the specified tolerance and constraint violation less than <code>options.TolCon</code>
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code>
-1	Stopped by an output function or plot function.
-2	No feasible point was found.

<code>lambda</code>	Structure containing the Lagrange multipliers at the solution <code>x</code> (separated by constraint type). The fields of the structure are
<code>lower</code>	Lower bounds <code>lb</code>
<code>upper</code>	Upper bounds <code>ub</code>
<code>ineqlin</code>	Linear inequalities
<code>eqlin</code>	Linear equalities
<code>ineqnonlin</code>	Nonlinear inequalities
<code>eqnonlin</code>	Nonlinear equalities
<code>output</code>	Structure containing information about the optimization. The fields of the structure are
<code>iterations</code>	Number of iterations taken
<code>funcCount</code>	Number of function evaluations
<code>lssteplength</code>	Size of final line search step relative to search direction
<code>stepsize</code>	Final displacement in <code>x</code>
<code>algorithm</code>	Optimization algorithm used
<code>firstorderopt</code>	Measure of first-order optimality
<code>constrviolation</code>	Maximum of constraint functions
<code>message</code>	Exit message

Options

Optimization options used by `fgoalattain`. Use `optimoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information.

DerivativeCheck	Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives. The choices are 'on' or the default, 'off'.
Diagnostics	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default, 'off'.
DiffMaxChange	Maximum change in variables for finite-difference gradients (a positive scalar). The default is Inf.
DiffMinChange	Minimum change in variables for finite-difference gradients (a positive scalar). The default is 0.
Display	Level of display. <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'iter' displays output at each iteration, and gives the default exit message.• 'iter-detailed' displays output at each iteration, and gives the technical exit message.• 'notify' displays output only if the function does not converge, and gives the default exit message.• 'notify-detailed' displays output only if the function does not converge, and gives the technical exit message.• 'final' (default) displays just the final output, and gives the default exit message.

FinDiffRelStep

- 'final-detailed' displays just the final output, and gives the technical exit message.

Scalar or vector step size factor. When you set FinDiffRelStep to a vector v , forward finite differences δ are

```
delta =  
v.*sign(x).*max(abs(x),TypicalX);
```

and central finite differences are

```
delta =  
v.*max(abs(x),TypicalX);
```

Scalar FinDiffRelStep expands to a vector. The default is $\sqrt{\text{eps}}$ for forward finite differences, and $\text{eps}^{(1/3)}$ for central finite differences.

FinDiffType

Finite differences, used to estimate gradients, are either 'forward' (default), or 'central' (centered). 'central' takes twice as many function evaluations, but should be more accurate.

The algorithm is careful to obey bounds when estimating both types of finite differences. So, for example, it could take a backward, rather than a forward, difference to avoid evaluating at a point outside bounds.

FunValCheck	Check whether objective function and constraints values are valid. 'on' displays an error when the objective function or constraints return a value that is complex, Inf, or NaN. The default, 'off', displays no error.
GoalsExactAchieve	Specifies the number of objectives for which it is required for the objective fun to equal the goal goal (a nonnegative integer). Such objectives should be partitioned into the first few elements of F. The default is 0.
GradConstr	Gradient for nonlinear constraint functions defined by the user. When set to 'on', fgoalattain expects the constraint function to have four outputs, as described in nonlcon in the “Input Arguments” on page 10-15 section. When set to the default, 'off', gradients of the nonlinear constraints are estimated by finite differences.
GradObj	Gradient for the objective function defined by the user. See the preceding description of fun to see how to define the gradient in fun. Set to 'on' to have fgoalattain use a user-defined gradient of the objective function. The default, 'off', causes fgoalattain to estimate gradients using finite differences.
MaxFunEvals	Maximum number of function evaluations allowed (a positive integer). The default is 100*numberOfVariables.

MaxIter	Maximum number of iterations allowed (a positive integer). The default is 400.
MaxSQPIter	Maximum number of SQP iterations allowed (a positive integer). The default is $10 \times \max(\text{numberOfVariables}, \text{numberOfInequalities} + \text{numberOfBounds})$
MeritFunction	Use goal attainment/minimax merit function if set to 'multiobj', the default. Use fmincon merit function if set to 'singleobj'.
OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none ([]). See “Output Function” on page 9-21.
PlotFcns	Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]). <ul style="list-style-type: none">• @optimplotx plots the current point• @optimplotfunccount plots the function count• @optimplotfval plots the function value• @optimplotconstrviolation plots the maximum constraint violation

	<ul style="list-style-type: none">• @optimplotstepsize plots the step size <p>For information on writing a custom plot function, see “Plot Functions” on page 9-30.</p>
RelLineSrchBnd	Relative bound (a real nonnegative scalar value) on the line search step length such that the total displacement in x satisfies $ \Delta x(i) \leq \text{relLineSrchBnd} \cdot \max(x(i) , typicalx(i))$. This option provides control over the magnitude of the displacements in x for cases in which the solver takes steps that are considered too large. The default is none ([]).
RelLineSrchBndDuration	Number of iterations for which the bound specified in <code>RelLineSrchBnd</code> should be active (default is 1).
TolCon	Termination tolerance on the constraint violation, a positive scalar. The default is $1e-6$.
TolConSQP	Termination tolerance on inner iteration SQP constraint violation, a positive scalar. The default is $1e-6$.
TolFun	Termination tolerance on the function value, a positive scalar. The default is $1e-6$.
TolX	Termination tolerance on x , a positive scalar. The default is $1e-6$.

TypicalX

Typical x values. The number of elements in TypicalX is equal to the number of elements in x0, the starting point. The default value is ones(numberofvariables,1). fgoalattain uses TypicalX for scaling finite differences for gradient estimation.

UseParallel

When true, estimate gradients in parallel. Disable by setting to the default, false. See “Parallel Computing”.

Examples

Consider a linear system of differential equations.

An output feedback controller, K, is designed producing a closed loop system

$$\begin{aligned}\dot{x} &= (A + BKC)x + Bu, \\ y &= Cx.\end{aligned}$$

The eigenvalues of the closed loop system are determined from the matrices A, B, C, and K using the command `eig(A+B*K*C)`. Closed loop eigenvalues must lie on the real axis in the complex plane to the left of the points [-5, -3, -1]. In order not to saturate the inputs, no element in K can be greater than 4 or be less than -4.

The system is a two-input, two-output, open loop, unstable system, with state-space matrices.

$$A = \begin{bmatrix} -0.5 & 0 & 0 \\ 0 & -2 & 10 \\ 0 & 1 & -2 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 \\ -2 & 2 \\ 0 & 1 \end{bmatrix} \quad C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The set of *goal values* for the closed loop eigenvalues is initialized as

$$\text{goal} = [-5, -3, -1];$$

To ensure the same percentage of under- or overattainment in the active objectives at the solution, the weighting matrix, `weight`, is set to `abs(goal)`.

Starting with a controller, $K = \begin{bmatrix} -1, -1; -1, -1 \end{bmatrix}$, first write a function file, `eigfun.m`.

```
function F = eigfun(K,A,B,C)
F = sort(eig(A+B*K*C)); % Evaluate objectives
```

Next, enter system matrices and invoke an optimization routine.

```
A = [-0.5 0 0; 0 -2 10; 0 1 -2];
B = [1 0; -2 2; 0 1];
C = [1 0 0; 0 0 1];
K0 = [-1 -1; -1 -1]; % Initialize controller matrix
goal = [-5 -3 -1]; % Set goal values for the eigenvalues
weight = abs(goal); % Set weight for same percentage
lb = -4*ones(size(K0)); % Set lower bounds on the controller
ub = 4*ones(size(K0)); % Set upper bounds on the controller
options = optimoptions('fgoalattain','Display','iter'); % Set display parameter
[K,fval,attainfactor] = fgoalattain(@(K)eigfun(K,A,B,C),...
    K0,goal,weight,[],[],[],[],lb,ub,[],options)
```

You can run this example by using the script `goaldemo`. (From the MATLAB Help browser or the MathWorks Web site documentation, you can click the `goaldemo` name to display the example.) After about 11 iterations, a solution is

Active inequalities (to within `options.TolCon = 1e-006`):

lower	upper	ineqlin	ineqnonlin
1			1
2			2
4			

```
K =
    -4.0000    -0.2564
```

```
-4.0000    -4.0000
```

```
fval =  
-6.9313  
-4.1588  
-1.4099
```

```
attainfactor =  
-0.3863
```

Discussion

The attainment factor indicates that each of the objectives has been overachieved by at least 38.63% over the original design goals. The active constraints, in this case constraints 1 and 2, are the objectives that are barriers to further improvement and for which the percentage of overattainment is met exactly. Three of the lower bound constraints are also active.

In the preceding design, the optimizer tries to make the objectives less than the goals. For a worst-case problem where the objectives must be as near to the goals as possible, use `optimoptions` to set the `GoalsExactAchieve` option to the number of objectives for which this is required.

Consider the preceding problem when you want all the eigenvalues to be equal to the goal values. A solution to this problem is found by invoking `fgoalattain` with the `GoalsExactAchieve` option set to 3.

```
options = optimoptions('fgoalattain','GoalsExactAchieve',3);  
[K,fval,attainfactor] = fgoalattain(...  
    @(K)eigfun(K,A,B,C),K0,goal,weight,[],[],[],[],lb,ub,[],...  
    options);
```

After about seven iterations, a solution is

```
K,fval,attainfactor
```

```
K =  
-1.5954    1.2040
```

```
-0.4201 -2.9046

fval =
-5.0000
-3.0000
-1.0000

attainfactor =
1.1304e-022
```

In this case the optimizer has tried to match the objectives to the goals. The attainment factor (of $1.1304e-22$ or so, depending on your system) indicates that the goals have been matched almost exactly.

Notes

This problem has discontinuities when the eigenvalues become complex; this explains why the convergence is slow. Although the underlying methods assume the functions are continuous, the method is able to make steps toward the solution because the discontinuities do not occur at the solution point. When the objectives and goals are complex, `fgoalattain` tries to achieve the goals in a least-squares sense.

Algorithms

Multiobjective optimization concerns the minimization of a set of objectives simultaneously. One formulation for this problem, and implemented in `fgoalattain`, is the goal attainment problem of Gembicki [3]. This entails the construction of a set of *goal* values for the objective functions. Multiobjective optimization is discussed in “Multiobjective Optimization Algorithms” on page 6-236.

In this implementation, the slack variable γ is used as a dummy argument to minimize the vector of objectives $F(x)$ simultaneously; *goal* is a set of values that the objectives attain. Generally, prior to the optimization, it is not known whether the objectives will reach the goals (under attainment) or be minimized less than the goals (overattainment). A weighting vector, *weight*, controls the relative underattainment or overattainment of the objectives.

`fgoalattain` uses a sequential quadratic programming (SQP) method, which is described in “Sequential Quadratic Programming (SQP)” on

page 6-33. Modifications are made to the line search and Hessian. In the line search an exact merit function (see [1] and [4]) is used together with the merit function proposed by [5] and [6]. The line search is terminated when either merit function shows improvement. A modified Hessian, which takes advantage of the special structure of the problem, is also used (see [1] and [4]). A full description of the modifications used is found in “Goal Attainment Method” on page 6-237 in “Introduction to Algorithms.” Setting the MeritFunction option to 'singleobj' with

```
options = optimoptions(options,'MeritFunction','singleobj')
```

uses the merit function and Hessian used in fmincon.

See also “SQP Implementation” on page 6-35 for more details on the algorithm used and the types of procedures displayed under the Procedures heading when the Display option is set to 'iter'.

Limitations

The objectives must be continuous. fgoalattain might give only local solutions.

References

- [1] Brayton, R.K., S.W. Director, G.D. Hachtel, and L.Vidigal, “A New Algorithm for Statistical Circuit Design Based on Quasi-Newton Methods and Function Splitting,” *IEEE Transactions on Circuits and Systems*, Vol. CAS-26, pp 784-794, Sept. 1979.
- [2] Fleming, P.J. and A.P. Pashkevich, *Computer Aided Control System Design Using a Multi-Objective Optimisation Approach*, Control 1985 Conference, Cambridge, UK, pp. 174-179.
- [3] Gembicki, F.W., “Vector Optimization for Control with Performance and Parameter Sensitivity Indices,” Ph.D. Dissertation, Case Western Reserve Univ., Cleveland, OH, 1974.
- [4] Grace, A.C.W., “Computer-Aided Control System Design Using Optimization Techniques,” Ph.D. Thesis, University of Wales, Bangor, Gwynedd, UK, 1989.

[5] Han, S.P., “A Globally Convergent Method For Nonlinear Programming,” *Journal of Optimization Theory and Applications*, Vol. 22, p. 297, 1977.

[6] Powell, M.J.D., “A Fast Algorithm for Nonlinear Constrained Optimization Calculations,” *Numerical Analysis*, ed. G.A. Watson, *Lecture Notes in Mathematics*, Vol. 630, Springer Verlag, 1978.

See Also

fmincon | fminimax | optimoptions | optimtool

How To

- `function_handle`
- “Multiobjective Optimization”

fminbnd

Purpose Find minimum of single-variable function on fixed interval

Equation Finds a minimum for a problem specified by

$$\min_x f(x) \text{ such that } x_1 < x < x_2.$$

x , x_1 , and x_2 are scalars and $f(x)$ is a function that returns a scalar.

Syntax

```
x = fminbnd(fun,x1,x2)
x = fminbnd(fun,x1,x2,options)
x = fminbnd(problem)
[x,fval] = fminbnd(...)
[x,fval,exitflag] = fminbnd(...)
[x,fval,exitflag,output] = fminbnd(...)
```

Description `fminbnd` attempts to find a minimum of a function of one variable within a fixed interval.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective function, if necessary.

`x = fminbnd(fun,x1,x2)` returns a value `x` that is a local minimizer of the scalar valued function that is described in `fun` in the interval $x_1 < x < x_2$. `fun` is either a function handle to a file or is an anonymous function.

`x = fminbnd(fun,x1,x2,options)` minimizes with the optimization options specified in the structure `options`. Use `optimset` to set these options.

`x = fminbnd(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-35.

Create the structure `problem` by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = fminbnd(...)` returns the value of the objective function computed in `fun` at the solution `x`.

`[x,fval,exitflag] = fminbnd(...)` returns a value `exitflag` that describes the exit condition of `fminbnd`.

`[x,fval,exitflag,output] = fminbnd(...)` returns a structure `output` that contains information about the optimization.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `fminbnd`. This section provides function-specific details for `fun`, `options`, and `problem`:

`fun` The function to be minimized. `fun` is a function handle for a function that accepts a scalar `x` and returns a scalar `f`, the objective function evaluated at `x`. The function `fun` can be specified as a function handle for a file:

```
x = fminbnd(@myfun,x1,x2)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...           % Compute function value at x.
```

`fun` can also be a function handle for an anonymous function.

```
x = fminbnd(@(x)sin(x^2),x1,x2);
```

`options` “Options” on page 10-37 provides the function-specific details for the options values.

fminbnd

problem	f	Objective function
	x1	Left endpoint
	x2	Right endpoint
solver	'fminbnd'	
options		Options structure created with optimset

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `fminbnd`. This section provides function-specific details for `exitflag` and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.
1	Function converged to a solution <code>x</code> .
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .
-1	Stopped by an output function or plot function.
-2	The bounds are inconsistent, meaning $x1 > x2$.
<code>output</code>	Structure containing information about the optimization. The fields of the structure are
<code>iterations</code>	Number of iterations taken
<code>funcCount</code>	Number of function evaluations
<code>algorithm</code>	Optimization algorithm used
<code>message</code>	Exit message

Options

Optimization options used by `fminbnd`. You can use `optimset` to set or change the values of these fields in the options structure `options`. See “Optimization Options Reference” on page 9-7 for detailed information.

<code>Display</code>	Level of display. 'off' or 'none' displays no output; 'iter' displays output at each iteration; 'final' displays just the final output; 'notify' (default) displays output only if the function does not converge.
<code>FunValCheck</code>	Check whether objective function values are valid. 'on' displays an error when the objective function returns a value that is complex, Inf, or NaN. The default 'off' displays no error.
<code>MaxFunEvals</code>	Maximum number of function evaluations allowed, a positive integer. The default is 500.
<code>MaxIter</code>	Maximum number of iterations allowed, a positive integer. The default is 500.
<code>OutputFcn</code>	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none (<code>[]</code>). See “Output Function” on page 9-21.
<code>PlotFcns</code>	Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none (<code>[]</code>). <ul style="list-style-type: none">• <code>@optimplotx</code> plots the current point• <code>@optimplotfunccount</code> plots the function count• <code>@optimplotfval</code> plots the function value

For information on writing a custom plot function, see “Plot Functions” on page 9-30.

`TolX` Termination tolerance on `x`, a positive scalar. The default is `1e-4`.

Examples

A minimum of $\sin(x)$ occurs at

```
x = fminbnd(@sin,0,2*pi)
x =
    4.7124
```

The value of the function at the minimum is

```
y = sin(x)
y =
   -1.0000
```

To find the minimum of the function

$$f(x) = (x - 3)^2 - 1,$$

on the interval $(0,5)$, first write a function file.

```
function f = myfun(x)
f = (x-3)^2 - 1;
```

Next, call an optimization routine.

```
x = fminbnd(@myfun,0,5)
```

This generates the solution

```
x =
    3
```

The value at the minimum is

```
y = myfun(x)
```

```
y =
    -1
```

If `fun` is parameterized, you can use anonymous functions to capture the problem-dependent parameters. For example, suppose you want to minimize the objective function `myfun` defined by the following function file:

```
function f = myfun(x,a)
f = (x - a)^2;
```

Note that `myfun` has an extra parameter `a`, so you cannot pass it directly to `fminbnd`. To optimize for a specific value of `a`, such as `a = 1.5`.

1 Assign the value to `a`.

```
a = 1.5; % define parameter first
```

2 Call `fminbnd` with a one-argument anonymous function that captures that value of `a` and calls `myfun` with two arguments:

```
x = fminbnd(@(x) myfun(x,a),0,1)
```

```
x =
    0.9999
```

Algorithms

`fminbnd` is a function file. The algorithm is based on golden section search and parabolic interpolation. Unless the left endpoint x_1 is very close to the right endpoint x_2 , `fminbnd` never evaluates `fun` at the endpoints, so `fun` need only be defined for x in the interval $x_1 < x < x_2$.

If the minimum actually occurs at x_1 or x_2 , `fminbnd` returns a point `x` in the interior of the interval (x_1, x_2) that is close to the minimizer. In this case, the distance of `x` from the minimizer is no more than $2*(TolX + 3*abs(x)*sqrt(eps))$. See [1] or [2] for details about the algorithm.

fminbnd

Limitations

The function to be minimized must be continuous. `fminbnd` might only give local solutions.

`fminbnd` often exhibits slow convergence when the solution is on a boundary of the interval. In such a case, `fmincon` often gives faster and more accurate solutions.

`fminbnd` only handles real variables.

References

[1] Forsythe, G.E., M.A. Malcolm, and C.B. Moler, *Computer Methods for Mathematical Computations*, Prentice Hall, 1976.

[2] Brent, Richard. P., *Algorithms for Minimization without Derivatives*, Prentice-Hall, Englewood Cliffs, New Jersey, 1973.

See Also

`fminsearch` | `fmincon` | `fminunc` | `optimset` | `optimtool`

How To

- `function_handle`
- “Anonymous Functions”

Purpose Find minimum of constrained nonlinear multivariable function

Equation Finds the minimum of a problem specified by

$$\min_x f(x) \text{ such that } \begin{cases} c(x) \leq 0 \\ ceq(x) = 0 \\ A \cdot x \leq b \\ Aeq \cdot x = beq \\ lb \leq x \leq ub, \end{cases}$$

b and beq are vectors, A and Aeq are matrices, $c(x)$ and $ceq(x)$ are functions that return vectors, and $f(x)$ is a function that returns a scalar. $f(x)$, $c(x)$, and $ceq(x)$ can be nonlinear functions.

x , lb , and ub can be passed as vectors or matrices; see “Matrix Arguments” on page 2-32.

Syntax

```
x = fmincon(fun,x0,A,b)
x = fmincon(fun,x0,A,b,Aeq,beq)
x = fmincon(fun,x0,A,b,Aeq,beq,lb,ub)
x = fmincon(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon)
x = fmincon(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)
x = fmincon(problem)
[x,fval] = fmincon(...)
[x,fval,exitflag] = fmincon(...)
[x,fval,exitflag,output] = fmincon(...)
[x,fval,exitflag,output,lambda] = fmincon(...)
[x,fval,exitflag,output,lambda,grad] = fmincon(...)
[x,fval,exitflag,output,lambda,grad,hessian] = fmincon(...)
```

Description

`fmincon` attempts to find a constrained minimum of a scalar function of several variables starting at an initial estimate. This is generally referred to as *constrained nonlinear optimization* or *nonlinear programming*.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective function and nonlinear constraint functions, if necessary.

$x = \text{fmincon}(\text{fun}, x_0, A, b)$ starts at x_0 and attempts to find a minimizer x of the function described in fun subject to the linear inequalities $A*x \leq b$. x_0 can be a scalar, vector, or matrix.

$x = \text{fmincon}(\text{fun}, x_0, A, b, Aeq, beq)$ minimizes fun subject to the linear equalities $Aeq*x = beq$ and $A*x \leq b$. If no inequalities exist, set $A = []$ and $b = []$.

$x = \text{fmincon}(\text{fun}, x_0, A, b, Aeq, beq, lb, ub)$ defines a set of lower and upper bounds on the design variables in x , so that the solution is always in the range $lb \leq x \leq ub$. If no equalities exist, set $Aeq = []$ and $beq = []$. If $x(i)$ is unbounded below, set $lb(i) = -\text{Inf}$, and if $x(i)$ is unbounded above, set $ub(i) = \text{Inf}$.

Note If the specified input bounds for a problem are inconsistent, the output x is x_0 and the output $fval$ is $[]$.

Components of x_0 that violate the bounds $lb \leq x \leq ub$ are reset to the interior of the box defined by the bounds. Components that respect the bounds are not changed.

See “Iterations Can Violate Constraints” on page 2-34.

$x = \text{fmincon}(\text{fun}, x_0, A, b, Aeq, beq, lb, ub, \text{nonlcon})$ subjects the minimization to the nonlinear inequalities $c(x)$ or equalities $ceq(x)$ defined in nonlcon . fmincon optimizes such that $c(x) \leq 0$ and $ceq(x) = 0$. If no bounds exist, set $lb = []$ and/or $ub = []$.

$x = \text{fmincon}(\text{fun}, x_0, A, b, Aeq, beq, lb, ub, \text{nonlcon}, \text{options})$ minimizes with the optimization options specified in options . Use

`optimoptions` to set these options. If there are no nonlinear inequality or equality constraints, set `nonlcon = []`.

`x = fmincon(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-43. Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = fmincon(...)` returns the value of the objective function `fun` at the solution `x`.

`[x,fval,exitflag] = fmincon(...)` returns a value `exitflag` that describes the exit condition of `fmincon`.

`[x,fval,exitflag,output] = fmincon(...)` returns a structure `output` with information about the optimization.

`[x,fval,exitflag,output,lambda] = fmincon(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

`[x,fval,exitflag,output,lambda,grad] = fmincon(...)` returns the value of the gradient of `fun` at the solution `x`.

`[x,fval,exitflag,output,lambda,grad,hessian] = fmincon(...)` returns the value of the Hessian at the solution `x`. See “fmincon Hessian” on page 3-29.

Input Arguments

“Function Arguments” on page 9-2 describes the arguments passed to `fmincon`. “Options” on page 10-52 provides the function-specific details for the options values. This section provides function-specific details for `fun`, `nonlcon`, and `problem`.

fun The function to be minimized. `fun` is a function that accepts a vector `x` and returns a scalar `f`, the objective function evaluated at `x`. `fun` can be specified as a function handle for a file:

```
x = fmincon(@myfun,x0,A,b)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...                    % Compute function value at x
```

`fun` can also be a function handle for an anonymous function:

```
x = fmincon(@(x)norm(x)^2,x0,A,b);
```

If the gradient of `fun` can also be computed *and* the `GradObj` option is `'on'`, as set by

```
options = optimoptions('fmincon','GradObj','on')
```

then `fun` must return the gradient vector $g(x)$ in the second output argument.

If the Hessian matrix can also be computed *and* the Hessian option is `'on'` via `options = optimoptions('fmincon','Hessian','user-supplied')` *and* the `Algorithm` option is `trust-region-reflective`, `fun` must return the Hessian value $H(x)$, a symmetric matrix, in a third output argument. `fun` can give a sparse Hessian. See “Writing Objective Functions” on page 2-18 for details.

If the Hessian matrix can be computed and the `Algorithm` option is `interior-point`, there are several ways to pass the Hessian to `fmincon`. For more information, see “Hessian” on page 10-49.

A, b, Aeq, beq Linear constraint matrices `A` and `Aeq`, and their corresponding vectors `b` and `beq`, can be sparse or dense. The `trust-region-reflective` and `interior-point` algorithms use sparse linear algebra. If `A` or `Aeq` is large, with relatively few nonzero entries, save running time and memory in the

trust-region-reflective or interior-point algorithms by using sparse matrices.

`nonlcon` The function that computes the nonlinear inequality constraints $c(x) \leq 0$ and the nonlinear equality constraints $ceq(x) = 0$. `nonlcon` accepts a vector x and returns the two vectors c and ceq . c is a vector that contains the nonlinear inequalities evaluated at x , and ceq is a vector that contains the nonlinear equalities evaluated at x . `nonlcon` should be specified as a function handle to a file or to an anonymous function, such as `mycon`:

```
x = fmincon(@myfun,x0,A,b,Aeq,beq,lb,ub,@mycon)
```

where `mycon` is a MATLAB function such as

```
function [c,ceq] = mycon(x)
c = ...      % Compute nonlinear inequalities at x.
ceq = ...    % Compute nonlinear equalities at x.
```

If the gradients of the constraints can also be computed *and* the `GradConstr` option is 'on', as set by

```
options = optimoptions('fmincon','GradConstr','on')
```

then `nonlcon` must also return, in the third and fourth output arguments, `GC`, the gradient of $c(x)$, and `GCEq`, the gradient of $ceq(x)$. `GC` and `GCEq` can be sparse or dense. If `GC` or `GCEq` is large, with relatively few nonzero entries, save running time and memory in the interior-point algorithm by representing them as sparse matrices. For more information, see “Nonlinear Constraints” on page 2-37.

Note Because Optimization Toolbox functions only accept inputs of type `double`, user-supplied objective and nonlinear constraint functions must return outputs of type `double`.

fmincon

problem	objective	Objective function
	x0	Initial point for x
	Aineq	Matrix for linear inequality constraints
	bineq	Vector for linear inequality constraints
	Aeq	Matrix for linear equality constraints
	beq	Vector for linear equality constraints
	lb	Vector of lower bounds
	ub	Vector of upper bounds
	nonlcon	Nonlinear constraint function
	solver	'fmincon'
	options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 describes arguments returned by `fmincon`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated. All Algorithms: 1 First-order optimality measure was less than <code>options.TolFun</code> , and maximum constraint violation was less than <code>options.TolCon</code> . 0 Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .
-----------------------	--

-1 Stopped by an output function or plot function.

-2 No feasible point was found.

trust-region-reflective, interior-point, and sqp algorithms:

2 Change in x was less than `options.TolX` and maximum constraint violation was less than `options.TolCon`.

trust-region-reflective algorithm only:

3 Change in the objective function value was less than `options.TolFun` and maximum constraint violation was less than `options.TolCon`.

active-set algorithm only:

4 Magnitude of the search direction was less than $2 * \text{options.TolX}$ and maximum constraint violation was less than `options.TolCon`.

5 Magnitude of directional derivative in search direction was less than $2 * \text{options.TolFun}$ and maximum constraint violation was less than `options.TolCon`.

interior-point and sqp algorithms:

-3 Objective function at current iteration went below `options.ObjectiveLimit` and maximum constraint violation was less than `options.TolCon`.

grad	Gradient at x
hessian	Hessian at x
lambda	Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields of the structure are: lower Lower bounds lb upper Upper bounds ub ineqlin Linear inequalities eqlin Linear equalities ineqnonlin Nonlinear inequalities eqnonlin Nonlinear equalities
output	Structure containing information about the optimization. The fields of the structure are: iterations Number of iterations taken funcCount Number of function evaluations lssteplength Size of line search step relative to search direction (active-set algorithm only) constrviolation Maximum of constraint functions stepsize Length of last displacement in x (active-set and interior-point algorithms) algorithm Optimization algorithm used cgiterations Total number of PCG iterations (trust-region-reflective and interior-point algorithms) firstorderopt Measure of first-order optimality message Exit message

Hessian

fmincon uses a Hessian as an optional input. This Hessian is the second derivatives of the Lagrangian (see Equation 3-1), namely,

$$\nabla_{xx}^2 L(x, \lambda) = \nabla^2 f(x) + \sum \lambda_i \nabla^2 c_i(x) + \sum \lambda_i \nabla^2 ceq_i(x). \quad (10-1)$$

The various fmincon algorithms handle input Hessians differently:

- The **active-set** and **sqp** algorithms do not accept a user-supplied Hessian. They compute a quasi-Newton approximation to the Hessian of the Lagrangian.
- The **trust-region-reflective** algorithm can accept a user-supplied Hessian as the final output of the objective function. Since this algorithm has only bounds or linear constraints, the Hessian of the Lagrangian is same as the Hessian of the objective function. See “Writing Scalar Objective Functions” on page 2-19 for details on how to pass the Hessian to fmincon. Indicate that you are supplying a Hessian by

```
options = optimoptions('fmincon','Algorithm','trust-region-reflective');
```

If you do not pass a Hessian, the algorithm computes a finite-difference approximation.

- The **interior-point** algorithm can accept a user-supplied Hessian as a separately defined function—it is not computed in the objective function. The syntax is

```
hessian = hessianfcn(x, lambda)
```

hessian is an n -by- n matrix, sparse or dense, where n is the number of variables. If **hessian** is large and has relatively few nonzero entries, save running time and memory by representing **hessian** as a sparse matrix. **lambda** is a structure with the Lagrange multiplier vectors associated with the nonlinear constraints:

```
lambda.ineqnonlin
```

`lambda.eqnonlin`

`fmincon` computes the structure `lambda`. `hessianfcn` must calculate the sums in Equation 10-1. Indicate that you are supplying a Hessian by

```
options = optimoptions('fmincon','Algorithm','interior-point',...  
    'Hessian','user-supplied','HessFcn',@hessianfcn);
```

For an example, see “`fmincon` Interior-Point Algorithm with Analytic Hessian” on page 6-57.

The interior-point algorithm has several more options for Hessians, see “Choose Input Hessian for interior-point `fmincon`” on page 2-23:

- `options = optimoptions('fmincon','Hessian','bfgs');`
`fmincon` calculates the Hessian by a dense quasi-Newton approximation. This is the default.
- `options = optimoptions('fmincon','Hessian','lbfgs');`
`fmincon` calculates the Hessian by a limited-memory, large-scale quasi-Newton approximation. The default memory, 10 iterations, is used.
- `options = optimoptions('fmincon','Hessian',{'lbfgs',positive integer});`
`fmincon` calculates the Hessian by a limited-memory, large-scale quasi-Newton approximation. The positive integer specifies how many past iterations should be remembered.
- `options = optimoptions('fmincon','Hessian','fin-diff-grads',...
 'SubproblemAlgorithm','cg','GradObj','on',...
 'GradConstr','on');`
`fmincon` calculates a Hessian-times-vector product by finite differences of the gradient(s). You must supply the gradient of the objective function, and also gradients of nonlinear constraints if they exist.

Hessian Multiply Function

The interior-point and trust-region-reflective algorithms allow you to supply a Hessian multiply function. This function gives the result of a Hessian-times-vector product, without computing the Hessian directly. This can save memory.

The syntax for the two algorithms differ:

- For the interior-point algorithm, the syntax is

```
W = HessMultFcn(x,lambda,v);
```

The result W should be the product $H*v$, where H is the Hessian of the Lagrangian at x (see Equation 10-1), λ is the Lagrange multiplier (computed by `fmincon`), and v is a vector of size n -by-1. Set options as follows:

```
options = optimoptions('fmincon','Algorithm','interior-point','HessMultFcn',  
    'SubproblemAlgorithm','cg','HessMult',@HessMultFcn);
```

Supply the function `HessMultFcn`, which returns an n -by-1 vector, where n is the number of dimensions of x . The `HessMult` option enables you to pass the result of multiplying the Hessian by a vector without calculating the Hessian.

- The trust-region-reflective algorithm does not involve λ :

```
W = HessMultFcn(H,v);
```

The result $W = H*v$. `fmincon` passes H as the value returned in the third output of the objective function (see “Writing Scalar Objective Functions” on page 2-19). `fmincon` also passes v , a vector or matrix with n rows. The number of columns in v can vary, so write `HessMultFcn` to accept an arbitrary number of columns. H does not have to be the Hessian; rather, it can be anything that enables you to calculate $W = H*v$.

Set options as follows:

```
options = optimoptions('fmincon','Algorithm','trust-region-reflective');
```

```
'Hessian','user-supplied','HessMult',@HessMultFcn);
```

For an example using a Hessian multiply function with the trust-region-reflective algorithm, see “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80.

Options

Optimization options used by `fmincon`. Some options apply to all algorithms, and others are relevant for particular algorithms. Use `optimoptions` to set or change the values in options. See “Optimization Options Reference” on page 9-7 for detailed information.

All Algorithms

All four algorithms use these options:

Algorithm Choose the optimization algorithm:

- 'interior-point' (default)
- 'trust-region-reflective'
- 'sqp'
- 'active-set'

For information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.

The trust-region-reflective algorithm requires:

- A gradient to be supplied in the objective function
- `GradObj` to be set to 'on'
- Either bound constraints or linear equality constraints, but not both

If you select the 'trust-region-reflective' algorithm and these conditions are not all satisfied, `fmincon` throws an error.

	The 'active-set' and 'sqp' algorithms are not large-scale. See “Large-Scale vs. Medium-Scale Algorithms” on page 2-12.
DerivativeCheck	Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives. The choices are 'on' or the default, 'off'.
Diagnostics	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default, 'off'.
DiffMaxChange	Maximum change in variables for finite-difference gradients (a positive scalar). The default is Inf.
DiffMinChange	Minimum change in variables for finite-difference gradients (a positive scalar). The default is 0.
Display	Level of display: <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'iter' displays output at each iteration, and gives the default exit message.• 'iter-detailed' displays output at each iteration, and gives the technical exit message.• 'notify' displays output only if the function does not converge, and gives the default exit message.• 'notify-detailed' displays output only if the function does not converge, and gives the technical exit message.• 'final' (default) displays just the final output, and gives the default exit message.

- 'final-detailed' displays just the final output, and gives the technical exit message.

`FinDiffRelStep` Scalar or vector step size factor. When you set `FinDiffRelStep` to a vector `v`, forward finite differences `delta` are

```
delta =  
v.*sign(x).*max(abs(x),TypicalX);
```

and central finite differences are

```
delta = v.*max(abs(x),TypicalX);
```

Scalar `FinDiffRelStep` expands to a vector. The default is `sqrt(eps)` for forward finite differences, and `eps^(1/3)` for central finite differences.

`FinDiffType` Finite differences, used to estimate gradients, are either 'forward' (default), or 'central' (centered). 'central' takes twice as many function evaluations but should be more accurate.

`fmincon` is careful to obey bounds when estimating both types of finite differences. So, for example, it could take a backward, rather than a forward, difference to avoid evaluating at a point outside bounds. However, for the interior-point algorithm, 'central' differences might violate bounds during their evaluation if the `AlwaysHonorConstraints` option is set to 'none'.

FunValCheck	Check whether objective function and constraints values are valid. 'on' displays an error when the objective function or constraints return a value that is complex, Inf, or NaN. The default, 'off', displays no error.
GradConstr	Gradient for nonlinear constraint functions defined by the user. When set to 'on', fmincon expects the constraint function to have four outputs, as described in nonlcon in the “Input Arguments” on page 10-43 section. When set to the default, 'off', gradients of the nonlinear constraints are estimated by finite differences. The trust-region-reflective algorithm does not accept nonlinear constraints.
GradObj	Gradient for the objective function defined by the user. See the preceding description of fun to see how to define the gradient in fun. Set to 'on' to have fmincon use a user-defined gradient of the objective function. The default, 'off', causes fmincon to estimate gradients using finite differences. You must provide the gradient, and set GradObj to 'on', to use the trust-region-reflective method.
MaxFunEvals	Maximum number of function evaluations allowed, a positive integer. The default value for all algorithms except interior-point is $100 \times \text{numberOfVariables}$; for the interior-point algorithm the default is 3000.
MaxIter	Maximum number of iterations allowed, a positive integer. The default value for all algorithms except interior-point is 400; for the interior-point algorithm the default is 1000.

OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none ([]). See “Output Function” on page 9-21.
PlotFcns	<p>Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]).</p> <ul style="list-style-type: none">• @optimplotx plots the current point• @optimplotfunccount plots the function count• @optimplotfval plots the function value• @optimplotconstrviolation plots the maximum constraint violation• @optimplotstepsize plots the step size• @optimplotfirstorderopt plots the first-order optimality measure <p>For information on writing a custom plot function, see “Plot Functions” on page 9-30.</p>
TolCon	Tolerance on the constraint violation, a positive scalar. The default is $1e-6$.
TolFun	Termination tolerance on the function value, a positive scalar. The default is $1e-6$.
TolX	Termination tolerance on x , a positive scalar. The default value for all algorithms except 'interior-point' is $1e-6$; for the 'interior-point' algorithm the default is $1e-10$.

TypicalX	<p>Typical x values. The number of elements in TypicalX is equal to the number of elements in x0, the starting point. The default value is ones(numberofvariables,1). fmincon uses TypicalX for scaling finite differences for gradient estimation.</p> <p>The 'trust-region-reflective' algorithm uses TypicalX only for the DerivativeCheck option.</p>
UseParallel	<p>When true, estimate gradients in parallel. Disable by setting to the default, false. trust-region-reflective requires a gradient in the objective, so UseParallel does not apply. See “Parallel Computing”.</p>

Trust-Region-Reflective Algorithm

The 'trust-region-reflective' algorithm uses these options:

Hessian	<p>If 'on' or 'user-supplied', fmincon uses a user-defined Hessian (defined in fun), or Hessian information (when using HessMult), for the objective function. If 'off' (default), fmincon approximates the Hessian using finite differences.</p>
HessMult	<p>Function handle for Hessian multiply function. For large-scale structured problems, this function computes the Hessian matrix product $H*Y$ without actually forming H. The function is of the form</p> $W = \text{hmfun}(\text{Hinfo}, Y)$ <p>where Hinfo contains a matrix used to compute $H*Y$.</p>

The first argument must be the same as the third argument returned by the objective function `fun`, for example:

```
[f,g,Hinfo] = fun(x)
```

`Y` is a matrix that has the same number of rows as there are dimensions in the problem. $W = H*Y$, although `H` is not formed explicitly. `fmincon` uses `Hinfo` to compute the preconditioner. See “Passing Extra Parameters” on page 2-53 for information on how to supply values for any additional parameters that `hmfun` needs.

Note Hessian must be set to 'on' or 'user-supplied' for `fmincon` to pass `Hinfo` from `fun` to `hmfun`.

See “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80 for an example.

HessPattern

Sparsity pattern of the Hessian for finite differencing. Set `HessPattern(i,j) = 1` when you can have $\partial^2 \text{fun} / \partial x(i) \partial x(j) \neq 0$. Otherwise, set `HessPattern(i,j) = 0`.

Use `HessPattern` when it is inconvenient to compute the Hessian matrix `H` in `fun`, but you can determine (say, by inspection) when the *i*th component of the gradient of `fun` depends on `x(j)`. `fmincon` can approximate `H` via sparse finite differences (of the gradient) if you provide the *sparsity structure* of `H` — i.e.,

locations of the nonzeros — as the value for `HessPattern`.

In the worst case, when the structure is unknown, do not set `HessPattern`. The default behavior is as if `HessPattern` is a dense matrix of ones. Then `fmincon` computes a full finite-difference approximation in each iteration. This can be very expensive for large problems, so it is usually better to determine the sparsity structure.

<code>MaxPCGIter</code>	Maximum number of PCG (preconditioned conjugate gradient) iterations, a positive scalar. The default is $\max(1, \text{floor}(\text{numberOfVariables}/2))$. For more information, see “Preconditioned Conjugate Gradient Method” on page 6-29.
<code>PrecondBandWidth</code>	Upper bandwidth of preconditioner for PCG, a nonnegative integer. By default, diagonal preconditioning is used (upper bandwidth of 0). For some problems, increasing the bandwidth reduces the number of PCG iterations. Setting <code>PrecondBandWidth</code> to <code>Inf</code> uses a direct factorization (Cholesky) rather than the conjugate gradients (CG). The direct factorization is computationally more expensive than CG, but produces a better quality step towards the solution.
<code>TolPCG</code>	Termination tolerance on the PCG iteration, a positive scalar. The default is 0.1.

Active-Set Algorithm

The 'active-set' algorithm uses these options:

MaxSQPIter	Maximum number of SQP iterations allowed, a positive integer. The default is $10 \cdot \max(\text{numberOfVariables}, \text{numberOfInequalities} + \text{numberOfBounds})$.
RelLineSrchBnd	Relative bound (a real nonnegative scalar value) on the line search step length such that the total displacement in x satisfies $ \Delta x(i) \leq \text{relLineSrchBnd} \cdot \max(x(i) , \text{typical}x(i))$. This option provides control over the magnitude of the displacements in x for cases in which the solver takes steps that are considered too large. The default is no bounds ([]).
RelLineSrchBndDuration	Number of iterations for which the bound specified in RelLineSrchBnd should be active (default is 1).
TolConSQP	Termination tolerance on inner iteration SQP constraint violation, a positive scalar. The default is $1e-6$.

Interior-Point Algorithm

The 'interior-point' algorithm uses these options:

AlwaysHonorConstraints	The default 'bounds' ensures that bound constraints are satisfied at every iteration. Disable by setting to 'none'.
HessFcn	Function handle to a user-supplied Hessian (see “Hessian” on page 10-49). This is used when the Hessian option is set to 'user-supplied'.
Hessian	Chooses how fmincon calculates the Hessian (see “Hessian” on page 10-49). The choices are: <ul style="list-style-type: none"> • 'bfgs' (default) • 'fin-diff-grads' • 'lbfgs' • {'lbfgs', Positive Integer} • 'user-supplied'
HessMult	Handle to a user-supplied function that gives a Hessian-times-vector product (see “Hessian” on page 10-49). This is used when the Hessian option is set to 'user-supplied'.
InitBarrierParam	Initial barrier value, a positive scalar. Sometimes it might help to try a value above the default 0.1, especially if the objective or constraint functions are large.
InitTrustRegionRadius	Initial radius of the trust region, a positive scalar. On badly scaled problems it might help to choose a value smaller than the default \sqrt{n} , where n is the number of variables.

MaxProjCGIter	A tolerance (stopping criterion) for the number of projected conjugate gradient iterations; this is an inner iteration, not the number of iterations of the algorithm. This positive integer has a default value of $2 * (\text{numberOfVariables} - \text{numberOfEqualities})$.
ObjectiveLimit	A tolerance (stopping criterion) that is a scalar. If the objective function value goes below <code>ObjectiveLimit</code> and the iterate is feasible, the iterations halt, since the problem is presumably unbounded. The default value is $-1 \text{e}20$.
ScaleProblem	'obj-and-constr' causes the algorithm to normalize all constraints and the objective function. Disable by setting to the default 'none'.
SubproblemAlgorithm	Determines how the iteration step is calculated. The default, 'ldl-factorization', is usually faster than 'cg' (conjugate gradient), though 'cg' might be faster for large problems with dense Hessians.
TolProjCG	A relative tolerance (stopping criterion) for projected conjugate gradient algorithm; this is for an inner iteration, not the algorithm iteration. This positive scalar has a default of 0.01.
TolProjCGAbs	Absolute tolerance (stopping criterion) for projected conjugate gradient algorithm; this is for an inner iteration, not the algorithm iteration. This positive scalar has a default of $1 \text{e}-10$.

SQP Algorithm

The 'sqp' algorithm uses these options:

ObjectiveLimit	A tolerance (stopping criterion) that is a scalar. If the objective function value goes below <code>ObjectiveLimit</code> and the iterate is feasible, the iterations halt, since the problem is presumably unbounded. The default value is <code>-1e20</code> .
ScaleProblem	'obj-and-constr' causes the algorithm to normalize all constraints and the objective function. Disable by setting to the default 'none'.

Examples

Find values of x that minimize $f(x) = -x_1x_2x_3$, starting at the point $x = [10; 10; 10]$, subject to the constraints:

$$0 \leq x_1 + 2x_2 + 2x_3 \leq 72.$$

- 1 Write a file that returns a scalar value f of the objective function evaluated at x :

```
function f = myfun(x)
f = -x(1) * x(2) * x(3);
```

- 2 Rewrite the constraints as both less than or equal to a constant,

$$\begin{aligned} -x_1 - 2x_2 - 2x_3 &\leq 0 \\ x_1 + 2x_2 + 2x_3 &\leq 72 \end{aligned}$$

- 3 Since both constraints are linear, formulate them as the matrix inequality $Ax \leq b$, where

```
A = [-1 -2 -2; ...
      1  2  2];
b = [0; 72];
```

4 Supply a starting point and invoke an optimization routine:

```
x0 = [10;10;10];    % Starting guess at the solution
[x,fval] = fmincon(@myfun,x0,A,b);
```

5 After fmincon stops, the solution is

```
x
x =
    24.0000
    12.0000
    12.0000
```

where the function value is

```
fval
fval =
   -3.4560e+03
```

and linear inequality constraints evaluate to be less than or equal to 0:

```
A*x - b
ans =
   -72.0000
    -0.0000
```

Notes

Trust-Region-Reflective Optimization

To use the trust-region-reflective algorithm, you must

- Supply the gradient of the objective function in `fun`.
- Set `GradObj` to `'on'` in `options`.
- Specify the feasible region using one, but not both, of the following types of constraints:
 - Upper and lower bounds constraints

- Linear equality constraints, in which the equality constraint matrix A_{eq} cannot have more rows than columns

You cannot use inequality constraints with the trust-region-reflective algorithm. If the preceding conditions are not met, `fmincon` reverts to the active-set algorithm.

`fmincon` returns a warning if you do not provide a gradient and the `Algorithm` option is 'trust-region-reflective'. `fmincon` permits an approximate gradient to be supplied, but this option is not recommended; the numerical behavior of most optimization methods is considerably more robust when the true gradient is used.

The trust-region-reflective method in `fmincon` is in general most effective when the matrix of second derivatives, i.e., the Hessian matrix $H(x)$, is also computed. However, evaluation of the true Hessian matrix is not required. For example, if you can supply the Hessian sparsity structure (using the `HessPattern` option in `options`), `fmincon` computes a sparse finite-difference approximation to $H(x)$.

If x_0 is not strictly feasible, `fmincon` chooses a new strictly feasible (centered) starting point.

If components of x have no upper (or lower) bounds, `fmincon` prefers that the corresponding components of `ub` (or `lb`) be set to `Inf` (or `-Inf` for `lb`) as opposed to an arbitrary but very large positive (or negative in the case of lower bounds) number.

Take note of these characteristics of linearly constrained minimization:

- A dense (or fairly dense) column of matrix A_{eq} can result in considerable fill and computational cost.
- `fmincon` removes (numerically) linearly dependent rows in A_{eq} ; however, this process involves repeated matrix factorizations and therefore can be costly if there are many dependencies.
- Each iteration involves a sparse least-squares solution with matrix

$$\overline{A_{eq}} = A_{eq}^T R^T,$$

where R^T is the Cholesky factor of the preconditioner. Therefore, there is a potential conflict between choosing an effective preconditioner and minimizing fill in \overline{Aeq} .

Active-Set Optimization

If equality constraints are present and dependent equalities are detected and removed in the quadratic subproblem, 'dependent' appears under the Procedures heading (when you ask for output by setting the Display option to 'iter'). The dependent equalities are only removed when the equalities are consistent. If the system of equalities is not consistent, the subproblem is infeasible and 'infeasible' appears under the Procedures heading.

Algorithms

Trust-Region-Reflective Optimization

The 'trust-region-reflective' algorithm is a subspace trust-region method and is based on the interior-reflective Newton method described in [3] and [4]. Each iteration involves the approximate solution of a large linear system using the method of preconditioned conjugate gradients (PCG). See the trust-region and preconditioned conjugate gradient method descriptions in “fmincon Trust Region Reflective Algorithm” on page 6-26.

Active-Set Optimization

fmincon uses a sequential quadratic programming (SQP) method. In this method, the function solves a quadratic programming (QP) subproblem at each iteration. fmincon updates an estimate of the Hessian of the Lagrangian at each iteration using the BFGS formula (see fminunc and references [7] and [8]).

fmincon performs a line search using a merit function similar to that proposed by [6], [7], and [8]. The QP subproblem is solved using an active set strategy similar to that described in [5]. “fmincon Active Set Algorithm” on page 6-32 describes this algorithm in detail.

See also “SQP Implementation” on page 6-35 for more details on the algorithm used.

Interior-Point Optimization

This algorithm is described in “fmincon Interior Point Algorithm” on page 6-43. There is more extensive description in [1], [41], and [9].

SQP Optimization

The fmincon 'sqp' algorithm is similar to the 'active-set' algorithm described in “Active-Set Optimization” on page 10-66. “fmincon SQP Algorithm” on page 6-42 describes the main differences. In summary, these differences are:

- “Strict Feasibility With Respect to Bounds” on page 6-42
- “Robustness to Non-Double Results” on page 6-42
- “Refactored Linear Algebra Routines” on page 6-42
- “Reformulated Feasibility Routines” on page 6-42

Limitations

fmincon is a gradient-based method that is designed to work on problems where the objective and constraint functions are both continuous and have continuous first derivatives.

When the problem is infeasible, fmincon attempts to minimize the maximum constraint value.

The 'trust-region-reflective' algorithm does not allow equal upper and lower bounds. For example, if $lb(2) == ub(2)$, fmincon gives this error:

```
Equal upper and lower bounds not permitted in this  
large-scale method.
```

```
Use equality constraints and the medium-scale  
method instead.
```

There are two different syntaxes for passing a Hessian, and there are two different syntaxes for passing a HessMult function; one for trust-region-reflective, and another for interior-point.

For `trust-region-reflective`, the Hessian of the Lagrangian is the same as the Hessian of the objective function. You pass that Hessian as the third output of the objective function.

For `interior-point`, the Hessian of the Lagrangian involves the Lagrange multipliers and the Hessians of the nonlinear constraint functions. You pass the Hessian as a separate function that takes into account both the position x and the Lagrange multiplier structure `lambda`.

Trust-Region-Reflective Coverage and Requirements

Additional Information Needed	For Large Problems
Must provide gradient for $f(x)$ in <code>fun</code> .	<ul style="list-style-type: none"> • Provide sparsity structure of the Hessian or compute the Hessian in <code>fun</code>. • The Hessian should be sparse. • <code>Aeq</code> should be sparse.

References

[1] Byrd, R.H., J. C. Gilbert, and J. Nocedal, “A Trust Region Method Based on Interior Point Techniques for Nonlinear Programming,” *Mathematical Programming*, Vol 89, No. 1, pp. 149–185, 2000.

[2] Byrd, R.H., Mary E. Hribar, and Jorge Nocedal, “An Interior Point Algorithm for Large-Scale Nonlinear Programming, SIAM Journal on Optimization,” *SIAM Journal on Optimization*, Vol 9, No. 4, pp. 877–900, 1999.

[3] Coleman, T.F. and Y. Li, “An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds,” *SIAM Journal on Optimization*, Vol. 6, pp. 418–445, 1996.

[4] Coleman, T.F. and Y. Li, “On the Convergence of Reflective Newton Methods for Large-Scale Nonlinear Minimization Subject to Bounds,” *Mathematical Programming*, Vol. 67, Number 2, pp. 189–224, 1994.

[5] Gill, P.E., W. Murray, and M.H. Wright, *Practical Optimization*, London, Academic Press, 1981.

[6] Han, S.P., “A Globally Convergent Method for Nonlinear Programming,” Vol. 22, *Journal of Optimization Theory and Applications*, p. 297, 1977.

[7] Powell, M.J.D., “A Fast Algorithm for Nonlinearly Constrained Optimization Calculations,” *Numerical Analysis*, ed. G.A. Watson, *Lecture Notes in Mathematics*, Springer Verlag, Vol. 630, 1978.

[8] Powell, M.J.D., “The Convergence of Variable Metric Methods For Nonlinearly Constrained Optimization Calculations,” *Nonlinear Programming 3* (O.L. Mangasarian, R.R. Meyer, and S.M. Robinson, eds.), Academic Press, 1978.

[9] Waltz, R. A., J. L. Morales, J. Nocedal, and D. Orban, “An interior algorithm for nonlinear optimization that combines line search and trust region steps,” *Mathematical Programming*, Vol 107, No. 3, pp. 391–408, 2006.

See Also

fminbnd | fminsearch | fminunc | optimoptions | optimtool

How To

- `function_handle`
- “Constrained Optimization”

fminimax

Purpose

Solve minimax constraint problem

Equation

Finds the minimum of a problem specified by

$$\min_x \max_i F_i(x) \text{ such that } \begin{cases} c(x) \leq 0 \\ ceq(x) = 0 \\ A \cdot x \leq b \\ Aeq \cdot x = beq \\ lb \leq x \leq ub \end{cases}$$

where b and beq are vectors, A and Aeq are matrices, and $c(x)$, $ceq(x)$, and $F(x)$ are functions that return vectors. $F(x)$, $c(x)$, and $ceq(x)$ can be nonlinear functions.

x , lb , and ub can be passed as vectors or matrices; see “Matrix Arguments” on page 2-32.

You can also solve max-min problems with `fminimax`, using the identity

$$\max_x \min_i F_i(x) = -\min_x \max_i (-F_i(x)).$$

You can solve problems of the form

$$\min_x \max_i |F_i(x)|$$

by using the `MinAbsMax` option; see “Notes” on page 10-83.

Syntax

```
x = fminimax(fun,x0)
x = fminimax(fun,x0,A,b)
x = fminimax(fun,x0,A,b,Aeq,beq)
x = fminimax(fun,x0,A,b,Aeq,beq,lb,ub)
x = fminimax(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon)
x = fminimax(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)
x = fminimax(problem)
[x,fval] = fminimax(...)
[x,fval,maxfval] = fminimax(...)
```

```
[x,fval,maxfval,exitflag] = fminimax(...)  
[x,fval,maxfval,exitflag,output] = fminimax(...)  
[x,fval,maxfval,exitflag,output,lambda] = fminimax(...)
```

Description

fminimax minimizes the worst-case (largest) value of a set of multivariable functions, starting at an initial estimate. This is generally referred to as the *minimax* problem.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective functions and nonlinear constraint functions, if necessary.

$x = \text{fminimax}(\text{fun},x0)$ starts at $x0$ and finds a minimax solution x to the functions described in fun .

$x = \text{fminimax}(\text{fun},x0,A,b)$ solves the minimax problem subject to the linear inequalities $A*x \leq b$.

$x = \text{fminimax}(\text{fun},x0,A,b,Aeq,beq)$ solves the minimax problem subject to the linear equalities $Aeq*x = beq$ as well. Set $A = []$ and $b = []$ if no inequalities exist.

$x = \text{fminimax}(\text{fun},x0,A,b,Aeq,beq,lb,ub)$ defines a set of lower and upper bounds on the design variables in x , so that the solution is always in the range $lb \leq x \leq ub$.

Note See “Iterations Can Violate Constraints” on page 2-34.

$x = \text{fminimax}(\text{fun},x0,A,b,Aeq,beq,lb,ub,\text{nonlcon})$ subjects the minimax problem to the nonlinear inequalities $c(x)$ or equality constraints $ceq(x)$ defined in nonlcon . fminimax optimizes such that $c(x) \leq 0$ and $ceq(x) = 0$. Set $lb = []$ and/or $ub = []$ if no bounds exist.

fminimax

`x = fminimax(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options.

`x = fminimax(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-72.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = fminimax(...)` returns the value of the objective function `fun` at the solution `x`.

`[x,fval,maxfval] = fminimax(...)` returns the maximum of the objective functions in the input `fun` evaluated at the solution `x`.

`[x,fval,maxfval,exitflag] = fminimax(...)` returns a value `exitflag` that describes the exit condition of `fminimax`.

`[x,fval,maxfval,exitflag,output] = fminimax(...)` returns a structure `output` with information about the optimization.

`[x,fval,maxfval,exitflag,output,lambda] = fminimax(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the output `fval` is `[]`.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `fminimax`. This section provides function-specific details for `fun`, `nonlcon`, and `problem`:

fun The function to be minimized. **fun** is a function that accepts a vector **x** and returns a vector **F**, the objective functions evaluated at **x**. The function **fun** can be specified as a function handle for a file:

```
x = fminimax(@myfun,x0)
```

where **myfun** is a MATLAB function such as

```
function F = myfun(x)
F = ...           % Compute function values at x
```

fun can also be a function handle for an anonymous function.

```
x = fminimax(@(x)sin(x.*x),x0);
```

If the user-defined values for **x** and **F** are matrices, they are converted to a vector using linear indexing.

To minimize the worst case absolute values of any of the elements of the vector $F(x)$ (i.e., $\min\{\max\{\text{abs}\{F(x)\}\}$), partition those objectives into the first elements of **F** and use **optimoptions** to set the **MinAbsMax** option to be the number of such objectives.

If the gradient of the objective function can also be computed *and* the **GradObj** option is 'on', as set by

```
options = optimoptions('fminimax','GradObj','on')
```

then the function **fun** must return, in the second output argument, the gradient value **G**, a matrix, at **x**. The gradient consists of the partial derivative dF/dx of each **F** at the point **x**. If **F** is a vector of length **m** and **x** has length **n**, where **n** is the length of **x0**, then the gradient **G** of **F**(**x**) is an **n**-by-**m** matrix where **G**(**i**, **j**) is the partial derivative of **F**(**j**) with respect to **x**(**i**) (i.e., the **j**th column of **G** is the gradient of the **j**th objective function **F**(**j**)).

By checking the value of **nargout**, the function can avoid computing **G** when **myfun** is called with only one output argument (in the case where the optimization algorithm only needs the value of **F** but not **G**).

```
function [F,G] = myfun(x)
F = ...           % Compute the function values at x
if nargin > 1     % Two output arguments
    G = ...       % Gradients evaluated at x
end
```

Note Setting GradObj to 'on' is effective only when there is no nonlinear constraint, or when the nonlinear constraint has GradConstr set to 'on' as well. This is because internally the objective is folded into the constraints, so the solver needs both gradients (objective and constraint) supplied in order to avoid estimating a gradient.

nonlcon The function that computes the nonlinear inequality constraints $c(x) \leq 0$ and nonlinear equality constraints $ceq(x) = 0$. The function `nonlcon` accepts a vector x and returns two vectors c and ceq . The vector c contains the nonlinear inequalities evaluated at x , and ceq contains the nonlinear equalities evaluated at x . The function `nonlcon` can be specified as a function handle.

```
x = fminimax(@myfun,x0,A,b,Aeq,beq,lb,ub,@mycon)
```

where `mycon` is a MATLAB function such as

```
function [c,ceq] = mycon(x)
c = ...           % Compute nonlinear inequalities at x
ceq = ...        % Compute nonlinear equalities at x
```

If the gradients of the constraints can also be computed *and* the GradConstr option is 'on', as set by

```
options = optimoptions('fminimax','GradConstr','on')
```

then the function `nonlcon` must also return, in the third and fourth output arguments, `GC`, the gradient of $c(x)$, and `GCeq`, the gradient of $ceq(x)$. “Nonlinear Constraints” on page 2-37 explains how to “conditionalize” the

gradients for use in solvers that do not accept supplied gradients, and explains the syntax of gradients.

Note Setting `GradConstr` to 'on' is effective only when `GradObj` is set to 'on' as well. This is because internally the objective is folded into the constraint, so the solver needs both gradients (objective and constraint) supplied in order to avoid estimating a gradient.

Note Because Optimization Toolbox functions only accept inputs of type `double`, user-supplied objective and nonlinear constraint functions must return outputs of type `double`.

<code>problem</code>	<code>objective</code>	Objective function
	<code>x0</code>	Initial point for <code>x</code>
	<code>Aineq</code>	Matrix for linear inequality constraints
	<code>bineq</code>	Vector for linear inequality constraints
	<code>Aeq</code>	Matrix for linear equality constraints
	<code>beq</code>	Vector for linear equality constraints
	<code>lb</code>	Vector of lower bounds
	<code>ub</code>	Vector of upper bounds
	<code>nonlcon</code>	Nonlinear constraint function
	<code>solver</code>	'fminimax'
	<code>options</code>	Options created with <code>optimoptions</code>

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `fminimax`. This section provides function-specific details for `exitflag`, `lambda`, `maxfval`, and `output`:

fminimax

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated:
1	Function converged to a solution <code>x</code> .
4	Magnitude of the search direction less than the specified tolerance and constraint violation less than <code>options.TolCon</code> .
5	Magnitude of directional derivative less than the specified tolerance and constraint violation less than <code>options.TolCon</code> .
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .
-1	Algorithm was terminated by the output function.
-2	No feasible point was found.
<code>lambda</code>	Structure containing the Lagrange multipliers at the solution <code>x</code> (separated by constraint type). The fields of the structure are
<code>lower</code>	Lower bounds <code>lb</code>
<code>upper</code>	Upper bounds <code>ub</code>
<code>ineqlin</code>	Linear inequalities
<code>eqlin</code>	Linear equalities
<code>ineqnonlin</code>	Nonlinear inequalities
<code>eqnonlin</code>	Nonlinear equalities
<code>maxfval</code>	Maximum of the function values evaluated at the solution <code>x</code> , that is, <code>maxfval = max{fun(x)}</code> .

output	Structure containing information about the optimization. The fields of the structure are	
	iterations	Number of iterations taken.
	funcCount	Number of function evaluations.
	lssteplength	Size of line search step relative to search direction
	stepsize	Final displacement in x
	algorithm	Optimization algorithm used.
	firstorderopt	Measure of first-order optimality
	constrviolation	Maximum of constraint functions
	message	Exit message

Options

Optimization options used by `fminimax`. Use `optimoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information.

DerivativeCheck	Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives. The choices are 'on' or the default, 'off'.
Diagnostics	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default, 'off'.
DiffMaxChange	Maximum change in variables for finite-difference gradients (a positive scalar). The default is Inf.
DiffMinChange	Minimum change in variables for finite-difference gradients (a positive scalar). The default is 0.

Display

Level of display:

- 'off' or 'none' displays no output.
- 'iter' displays output at each iteration, and gives the default exit message.
- 'iter-detailed' displays output at each iteration, and gives the technical exit message.
- 'notify' displays output only if the function does not converge, and gives the default exit message.
- 'notify-detailed' displays output only if the function does not converge, and gives the technical exit message.
- 'final' (default) displays just the final output, and gives the default exit message.
- 'final-detailed' displays just the final output, and gives the technical exit message.

FinDiffRelStep

Scalar or vector step size factor. When you set `FinDiffRelStep` to a vector `v`, forward finite differences `delta` are

```
delta =  
v.*sign(x).*max(abs(x),TypicalX);
```

and central finite differences are

```
delta =  
v.*max(abs(x),TypicalX);
```

	Scalar <code>FinDiffRelStep</code> expands to a vector. The default is <code>sqrt(eps)</code> for forward finite differences, and <code>eps^(1/3)</code> for central finite differences.
<code>FinDiffType</code>	Finite differences, used to estimate gradients, are either <code>'forward'</code> (the default), or <code>'central'</code> (centered). <code>'central'</code> takes twice as many function evaluations, but should be more accurate. The algorithm is careful to obey bounds when estimating both types of finite differences. So, for example, it could take a backward, rather than a forward, difference to avoid evaluating at a point outside bounds.
<code>FunValCheck</code>	Check whether objective function and constraints values are valid. <code>'on'</code> displays an error when the objective function or constraints return a value that is complex, <code>Inf</code> , or <code>NaN</code> . The default <code>'off'</code> displays no error.
<code>GradConstr</code>	Gradient for the user-defined constraints. When set to <code>'on'</code> , <code>fminimax</code> expects the constraint function to have four outputs, as described in <code>nonlcon</code> in “Input Arguments” on page 10-72. When set to the default <code>'off'</code> , <code>fminimax</code> estimates gradients of the nonlinear constraints by finite differences.

GradObj	Gradient for the user-defined objective function. See the preceding description of fun to see how to define the gradient in fun. Set to 'on' to have fminimax use a user-defined gradient of the objective function. The default 'off' causes fminimax to estimate gradients using finite differences.
MaxFunEvals	Maximum number of function evaluations allowed, a positive integer. The default value is $100 \times \text{numberOfVariables}$.
MaxIter	Maximum number of iterations allowed, a positive integer. The default value is 400.
MaxSQPIter	Maximum number of SQP iterations allowed, a positive integer. The default is $10 \times \max(\text{numberOfVariables}, \text{numberOfInequalities} + \text{numberOfBounds})$.
MeritFunction	Use the goal attainment/minimax merit function if set to 'multiobj' (default). Use the fmincon merit function if set to 'singleobj'.
MinAbsMax	Number of elements of $F_i(x)$ to minimize the maximum absolute value of F_i . See "Notes" on page 10-83. The default is 0.
OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none ([]). See "Output Function" on page 9-21.

PlotFcns

Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]).

- @optimplotx plots the current point.
- @optimplotfunccount plots the function count.
- @optimplotfval plots the function value.
- @optimplotconstrviolation plots the maximum constraint violation.
- @optimplotstepsize plots the step size.

For information on writing a custom plot function, see “Plot Functions” on page 9-30.

RelLineSrchBnd

Relative bound (a real nonnegative scalar value) on the line search step length such that the total displacement in x satisfies $|\Delta x(i)| \leq \text{relLineSrchBnd} \cdot \max(|x(i)|, |typicalx(i)|)$. This option provides control over the magnitude of the displacements in x for cases in which the solver takes steps that it considers too large. The default is no bounds ([]).

RelLineSrchBndDuration

Number of iterations for which the bound specified in **RelLineSrchBnd** should be active (default is 1).

fminimax

TolCon	Termination tolerance on the constraint violation, a positive scalar. The default is $1e-6$.
TolConSQP	Termination tolerance on inner iteration SQP constraint violation, a positive scalar. The default is $1e-6$.
TolFun	Termination tolerance on the function value, a positive scalar. The default is $1e-6$.
TolX	Termination tolerance on x , a positive scalar. The default value is $1e-6$.
TypicalX	Typical x values. The number of elements in <code>TypicalX</code> is equal to the number of elements in <code>x0</code> , the starting point. The default value is <code>ones(numberofvariables,1)</code> . <code>fminimax</code> uses <code>TypicalX</code> for scaling finite differences for gradient estimation.
UseParallel	When true, estimate gradients in parallel. Disable by setting to the default false. See “Parallel Computing”.

Examples

Find values of x that minimize the maximum value of

$$[f_1(x), f_2(x), f_3(x), f_4(x), f_5(x)]$$

where

$$f_1(x) = 2x_1^2 + x_2^2 - 48x_1 - 40x_2 + 304,$$

$$f_2(x) = -x_1^2 - 3x_2^2,$$

$$f_3(x) = x_1 + 3x_2 - 18,$$

$$f_4(x) = -x_1 - x_2,$$

$$f_5(x) = x_1 + x_2 - 8.$$

First, write a file that computes the five functions at x .

```
function f = myfun(x)
f(1) = 2*x(1)^2+x(2)^2-48*x(1)-40*x(2)+304;    % Objectives
f(2) = -x(1)^2 - 3*x(2)^2;
f(3) = x(1) + 3*x(2) - 18;
f(4) = -x(1) - x(2);
f(5) = x(1) + x(2) - 8;
```

Next, invoke an optimization routine.

```
x0 = [0.1; 0.1];    % Make a starting guess at solution
[x,fval] = fminimax(@myfun,x0);
```

After seven iterations, the solution is

```
x,fval
x =
    4.0000
    4.0000
fval =
    0.0000   -64.0000   -2.0000   -8.0000   -0.0000
```

Notes

You can solve problems of the form

$$\min_x \max_i G_i(x),$$

where

$$G_i(x) = \begin{cases} |F_i(x)| & 1 \leq i \leq m \\ F_i(x) & i > m. \end{cases}$$

Here m is the value of the `MinAbsMax` option. The advantage of this formulation is you can minimize the absolute value of some components of F , even though the absolute value function is not smooth.

In order to use this option, reorder the elements of F , if necessary, so the first elements are those for which you want the minimum absolute value.

For example, consider the problem in “Examples” on page 10-82. Modify the problem to find the minimum of the maximum absolute values of all $f_i(x)$. Solve this problem by invoking `fminimax` with the commands

```
x0 = [0.1; 0.1]; % Make a starting guess at the solution
options = optimoptions('fminimax','MinAbsMax',5); % Minimize abs. values
[x,fval] = fminimax(@myfun,x0,...
                  [],[],[],[],[],[],[],[],options);
```

After seven iterations, the solution is

```
x =
    4.9256
    2.0796
fval =
    37.2356 -37.2356 -6.8357 -7.0052 -0.9948
```

Algorithms

`fminimax` internally reformulates the minimax problem into an equivalent Nonlinear Linear Programming problem by appending additional (reformulation) constraints of the form $F_i(x) \leq \gamma$ to the constraints given in “Equation” on page 10-70, and then minimizing γ over x . `fminimax` uses a sequential quadratic programming (SQP) method [1] to solve this problem.

Modifications are made to the line search and Hessian. In the line search an exact merit function (see [2] and [4]) is used together with the merit function proposed by [3] and [5]. The line search is terminated

when either merit function shows improvement. The function uses a modified Hessian that takes advantage of the special structure of this problem. Using `optimoptions` to set the `MeritFunction` option to `'singleobj'` uses the merit function and Hessian used in `fmincon`.

See also “SQP Implementation” on page 6-35 for more details on the algorithm used and the types of procedures printed under the `Procedures` heading when you set the `Display` option to `'iter'`.

Limitations

The function to be minimized must be continuous. `fminimax` might only give local solutions.

References

[1] Brayton, R.K., S.W. Director, G.D. Hachtel, and L.Vidigal, “A New Algorithm for Statistical Circuit Design Based on Quasi-Newton Methods and Function Splitting,” *IEEE Trans. Circuits and Systems*, Vol. CAS-26, pp. 784-794, Sept. 1979.

[2] Grace, A.C.W., “Computer-Aided Control System Design Using Optimization Techniques,” Ph.D. Thesis, University of Wales, Bangor, Gwynedd, UK, 1989.

[3] Han, S.P., “A Globally Convergent Method For Nonlinear Programming,” *Journal of Optimization Theory and Applications*, Vol. 22, p. 297, 1977.

[4] Madsen, K. and H. Schjaer-Jacobsen, “Algorithms for Worst Case Tolerance Optimization,” *IEEE Trans. of Circuits and Systems*, Vol. CAS-26, Sept. 1979.

[5] Powell, M.J.D., “A Fast Algorithm for Nonlinearly Constrained Optimization Calculations,” *Numerical Analysis*, ed. G.A. Watson, *Lecture Notes in Mathematics*, Vol. 630, Springer Verlag, 1978.

See Also

`fgoalattain` | `lsqnonlin` | `optimoptions` | `optimtool`

How To

• `function_handle`

- “Multiobjective Optimization”

Purpose Find minimum of unconstrained multivariable function using derivative-free method

Equation Finds the minimum of a problem specified by

$$\min_x f(x)$$

where $f(x)$ is a function that returns a scalar.

x is a vector or a matrix; see “Matrix Arguments” on page 2-32.

Syntax

```
x = fminsearch(fun,x0)
x = fminsearch(fun,x0,options)
x = fminsearch(problem)
[x,fval] = fminsearch(...)
[x,fval,exitflag] = fminsearch(...)
[x,fval,exitflag,output] = fminsearch(...)
```

Description `fminsearch` attempts to find a minimum of a scalar function of several variables, starting at an initial estimate. This is generally referred to as *unconstrained nonlinear optimization*.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective function, if necessary.

`x = fminsearch(fun,x0)` starts at the point `x0` and returns a value `x` that is a local minimizer of the function described in `fun`. `fun` is either a function handle to a file or is an anonymous function. `x0` can be a scalar, vector, or matrix.

`x = fminsearch(fun,x0,options)` minimizes with the optimization options specified in the structure `options`. Use `optimset` to set these options.

`x = fminsearch(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-88.

fminsearch

Create the structure `problem` by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = fminsearch(...)` returns in `fval` the value of the objective function `fun` at the solution `x`.

`[x,fval,exitflag] = fminsearch(...)` returns a value `exitflag` that describes the exit condition of `fminsearch`.

`[x,fval,exitflag,output] = fminsearch(...)` returns a structure `output` that contains information about the optimization.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `fminsearch`. This section provides function-specific details for `fun`, `options`, and `problem`:

fun The function to be minimized. `fun` is a function handle for a function that accepts a vector `x` and returns a scalar `f`, the objective function evaluated at `x`. The function `fun` can be specified as a function handle for a file:

```
x = fminsearch(@myfun,x0)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...                    % Compute function value at x
```

`fun` can also be a function handle for an anonymous function, such as

```
x = fminsearch(@(x)norm(x)^2,x0,A,b);
```

options “Options” on page 10-90 provides the function-specific details for the options values.

problem	objective	Objective function
	x0	Initial point for x
	solver	'fminsearch'
	options	Options structure created with optimset

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `fminsearch`. This section provides function-specific details for `exitflag` and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.
1	The function converged to a solution <code>x</code> .
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .
-1	The algorithm was terminated by the output function.
<code>output</code>	Structure containing information about the optimization. The fields of the structure are
<code>iterations</code>	Number of iterations
<code>funcCount</code>	Number of function evaluations
<code>algorithm</code>	'Nelder-Mead simplex direct search'
<code>message</code>	Exit message

Options

Optimization options used by `fminsearch`. You can use `optimset` to set or change the values of these fields in the options structure `options`. See “Optimization Options Reference” on page 9-7 for detailed information.

<code>Display</code>	Level of display: <ul style="list-style-type: none">• <code>'off'</code> or <code>'none'</code> displays no output.• <code>'iter'</code> displays output at each iteration.• <code>'notify'</code> displays output only if the function does not converge.• <code>'final'</code> (default) displays just the final output.
<code>FunValCheck</code>	Check whether objective function values are valid. <code>'on'</code> displays an error when the objective function returns a value that is complex, <code>Inf</code> , or <code>NaN</code> . The default <code>'off'</code> displays no error.
<code>MaxFunEvals</code>	Maximum number of function evaluations allowed, a positive integer. The default is <code>200*numberOfVariables</code> .
<code>MaxIter</code>	Maximum number of iterations allowed, a positive integer. The default value is <code>200*numberOfVariables</code> .
<code>OutputFcn</code>	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is <code>none ([])</code> . See “Output Function” on page 9-21.

PlotFcns	<p>Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]):</p> <ul style="list-style-type: none"> • @optimplotx plots the current point. • @optimplotfunccount plots the function count. • @optimplotfval plots the function value. <p>For information on writing a custom plot function, see “Plot Functions” on page 9-30.</p>
TolFun	Termination tolerance on the function value, a positive scalar. The default is 1e-4.
TolX	Termination tolerance on x, a positive scalar. The default value is 1e-4.

Examples

Example 1: Minimizing Rosenbrock’s Function with fminsearch

A classic test example for multidimensional minimization is the Rosenbrock banana function:

$$f(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2.$$

The minimum is at (1, 1) and has the value 0. The traditional starting point is (-1.2, 1). The anonymous function shown here defines the function and returns a function handle called banana:

```
banana = @(x)100*(x(2)-x(1)^2)^2+(1-x(1))^2;
```

Pass the function handle to fminsearch:

```
[x,fval,exitflag] = fminsearch(banana,[-1.2, 1])
```

This produces

```
x =  
    1.0000    1.0000
```

```
fval =  
    8.1777e-010
```

```
exitflag =  
         1
```

This indicates that the minimizer was found at [1 1] with a value near zero.

Example 2

You can modify the first example by adding a parameter a to the second term of the banana function:

$$f(x) = 100(x_2 - x_1^2)^2 + (a - x_1)^2.$$

This changes the location of the minimum to the point $[a, a^2]$. To minimize this function for a specific value of a , for example $a = \sqrt{2}$, create a one-argument anonymous function that captures the value of a .

```
a = sqrt(2);  
banana = @(x) 100*(x(2)-x(1)^2)^2+(a-x(1))^2;
```

Then the statement

```
[x,fval,exitflag] = fminsearch(banana, [-1.2, 1], ...  
    optimset('TolX',1e-8))
```

seeks the minimum $[\sqrt{2}, 2]$ to an accuracy higher than the default on x . The result is

```
x =  
    1.4142    2.0000
```

```
fval =  
    4.2065e-018  
  
exitflag =  
    1
```

Algorithms

fminsearch uses the simplex search method of [1]. This is a direct search method that does not use numerical or analytic gradients as in fminunc. The algorithm is described in detail in “fminsearch Algorithm” on page 6-14.

fminsearch is generally less efficient than fminunc for problems of order greater than two. However, when the problem is highly discontinuous, fminsearch might be more robust.

Limitations

fminsearch solves nondifferentiable problems and can often handle discontinuity, particularly if it does not occur near the solution. fminsearch might only give local solutions.

fminsearch only minimizes over the real numbers, that is, x must only consist of real numbers and $f(x)$ must only return real numbers. When x has complex variables, they must be split into real and imaginary parts.

Notes

fminsearch is not the preferred choice for solving problems that are sums of squares, that is, of the form

$$\min_x \|f(x)\|_2^2 = \min_x (f_1(x)^2 + f_2(x)^2 + \dots + f_n(x)^2)$$

Instead use the lsqnonlin function, which has been optimized for problems of this form.

References

[1] Lagarias, J. C., J. A. Reeds, M. H. Wright, and P. E. Wright, “Convergence Properties of the Nelder-Mead Simplex Method in Low Dimensions,” *SIAM Journal of Optimization*, Vol. 9, Number 1, pp. 112–147, 1998.

fminsearch

See Also

fminbnd | fminunc | optimset | optimtool

How To

- `function_handle`
- “Anonymous Functions”

Purpose

Find minimum of unconstrained multivariable function

Equation

Finds the minimum of a problem specified by

$$\min_x f(x)$$

where $f(x)$ is a function that returns a scalar.

x is a vector or a matrix; see “Matrix Arguments” on page 2-32.

Syntax

```
x = fminunc(fun,x0)
x = fminunc(fun,x0,options)
x = fminunc(problem)
[x,fval] = fminunc(...)
[x,fval,exitflag] = fminunc(...)
[x,fval,exitflag,output] = fminunc(...)
[x,fval,exitflag,output,grad] = fminunc(...)
[x,fval,exitflag,output,grad,hessian] = fminunc(...)
```

Description

fminunc attempts to find a minimum of a scalar function of several variables, starting at an initial estimate. This is generally referred to as *unconstrained nonlinear optimization*.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective function, if necessary.

`x = fminunc(fun,x0)` starts at the point `x0` and attempts to find a local minimum `x` of the function described in `fun`. `x0` can be a scalar, vector, or matrix.

`x = fminunc(fun,x0,options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options.

`x = fminunc(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-96.

fminunc

Create the problem structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = fminunc(...)` returns in `fval` the value of the objective function `fun` at the solution `x`.

`[x,fval,exitflag] = fminunc(...)` returns a value `exitflag` that describes the exit condition.

`[x,fval,exitflag,output] = fminunc(...)` returns a structure `output` that contains information about the optimization.

`[x,fval,exitflag,output,grad] = fminunc(...)` returns in `grad` the value of the gradient of `fun` at the solution `x`.

`[x,fval,exitflag,output,grad,hessian] = fminunc(...)` returns in `hessian` the value of the Hessian of the objective function `fun` at the solution `x`. See “Hessian” on page 10-99.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `fminunc`. This section provides function-specific details for `fun`, `options`, and `problem`:

`fun`

The function to be minimized. `fun` is a function that accepts a vector `x` and returns a scalar `f`, the objective function evaluated at `x`. The function `fun` can be specified as a function handle for a file

```
x = fminunc(@myfun,x0)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...           % Compute function value at x
```

`fun` can also be a function handle for an anonymous function.

```
x = fminunc(@(x)norm(x)^2,x0);
```

If the gradient of `fun` can also be computed *and* the `GradObj` option is `'on'`, as set by

```
options = optimoptions(@fminunc,'GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `g`, a vector, at `x`. The gradient is the partial derivatives $\partial f / \partial x_i$ of `f` at the point `x`. That is, the *i*th component of `g` is the partial derivative of `f` with respect to the *i*th component of `x`.

If the Hessian matrix can also be computed *and* the Hessian option is `'on'`, i.e., `options = optimoptions(@fminunc,'GradObj','on','Hessian','on')`, then the function `fun` must return the Hessian value `H`, a symmetric matrix, at `x` in a third output argument. The Hessian matrix is the second partial derivatives matrix of `f` at the point `x`. That is, the (*i,j*)th component of `H` is the second partial derivative of `f` with respect to x_i and x_j , $\partial^2 f / \partial x_i \partial x_j$. The Hessian is by definition a symmetric matrix.

“Writing Scalar Objective Functions” on page 2-19 explains how to “conditionalize” the gradients and Hessians for use in solvers that do not accept them. “Passing Extra Parameters” on page 2-53 explains how to parameterize `fun`, if necessary.

options	“Options” on page 10-99 provides the function-specific details for the options values.	
problem	objective	Objective function
	x0	Initial point for x
	solver	'fminunc'
	options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `fminunc`. This section provides function-specific details for `exitflag` and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.
	1 Magnitude of gradient smaller than the <code>TolFun</code> tolerance.
	2 Change in x was smaller than the <code>TolX</code> tolerance.
	3 Change in the objective function value was less than the <code>TolFun</code> tolerance.
	5 Predicted decrease in the objective function was less than the <code>TolFun</code> tolerance.
	0 Number of iterations exceeded <code>MaxIter</code> or number of function evaluations exceeded <code>MaxFunEvals</code> .
	-1 Algorithm was terminated by the output function.
	-3 Objective function at current iteration went below <code>ObjectiveLimit</code> .
<code>grad</code>	Gradient at x
<code>hessian</code>	Hessian at x

<code>output</code>	Structure containing information about the optimization. The fields of the structure are
<code>iterations</code>	Number of iterations taken
<code>funcCount</code>	Number of function evaluations
<code>firstorderopt</code>	Measure of first-order optimality
<code>algorithm</code>	Optimization algorithm used
<code>cgiterations</code>	Total number of PCG iterations (trust-region algorithm only)
<code>stepsize</code>	Final displacement in x (medium-scale algorithm only)
<code>message</code>	Exit message

Hessian

`fminunc` computes the output argument `hessian` as follows:

- When using the medium-scale algorithm, the function computes a finite-difference approximation to the Hessian at x using
 - The gradient `grad` if you supply it
 - The objective function `fun` if you do not supply the gradient
- When using the trust-region algorithm, the function uses
 - `options.Hessian`, if you supply it, to compute the Hessian at x
 - A finite-difference approximation to the Hessian at x , if you supply only the gradient

Options

`fminunc` uses these optimization options. Some options apply to all algorithms, some are only relevant when you are using the trust-region algorithm, and others are only relevant when you are using the quasi-newton algorithm. Use `optimoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information.

All Algorithms

All `fminunc` algorithms use the following options:

Algorithm	Choose the <code>fminunc</code> algorithm. Choices are 'quasi-newton' and 'trust-region' (default).
If you use <code>optimset</code> (not recommended), use <code>LargeScale</code> instead of <code>Algorithm</code> .	The trust-region algorithm requires you to provide the gradient (see the preceding description of <code>fun</code>), or else <code>fminunc</code> uses the 'quasi-newton' algorithm. For information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.
<code>DerivativeCheck</code>	Compare user-supplied derivatives (gradient of objective) to finite-differencing derivatives. The choices are 'on' or the default 'off'.
<code>Diagnostics</code>	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default 'off'.
<code>DiffMaxChange</code>	Maximum change in variables for finite-difference gradients (a positive scalar). The default is <code>Inf</code> .
<code>DiffMinChange</code>	Minimum change in variables for finite-difference gradients (a positive scalar). The default is 0.
<code>Display</code>	Level of display: <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'iter' displays output at each iteration, and gives the default exit message.• 'iter-detailed' displays output at each iteration, and gives the technical exit message.• 'notify' displays output only if the function does not converge, and gives the default exit message.

- 'notify-detailed' displays output only if the function does not converge, and gives the technical exit message.
- 'final' (default) displays just the final output, and gives the default exit message.
- 'final-detailed' displays just the final output, and gives the technical exit message.

FinDiffRelStep Scalar or vector step size factor. When you set **FinDiffRelStep** to a vector v , forward finite differences δ are

$$\delta = v \cdot \text{sign}(x) \cdot \max(\text{abs}(x), \text{TypicalX});$$

and central finite differences are

$$\delta = v \cdot \max(\text{abs}(x), \text{TypicalX});$$

Scalar **FinDiffRelStep** expands to a vector. The default is $\sqrt{\text{eps}}$ for forward finite differences, and $\text{eps}^{(1/3)}$ for central finite differences.

The trust-region algorithm uses **FinDiffRelStep** only when **DerivativeCheck** is 'on'.

FinDiffType Finite differences, used to estimate gradients, are either 'forward' (the default), or 'central' (centered). 'central' takes twice as many function evaluations, but should be more accurate. The trust-region algorithm uses **FinDiffType** only when **DerivativeCheck** is 'on'.

FunValCheck	Check whether objective function values are valid. 'on' displays an error when the objective function returns a value that is complex, Inf, or NaN. The default, 'off', displays no error.
GradObj	Gradient for the objective function defined by the user. See the preceding description of fun to see how to define the gradient in fun. Set to 'on' to have fminunc use a user-defined gradient of the objective function. The default 'off' causes fminunc to estimate gradients using finite differences. You must provide the gradient, and set GradObj to 'on', to use the trust-region algorithm. This option is not required for the quasi-Newton algorithm.
LargeScale	Use large-scale algorithm if possible when set to the default 'on'. Use medium-scale algorithm when set to 'off'.
If you use optimoptions (recommended), use Algorithm instead of LargeScale.	The LargeScale algorithm requires you to provide the gradient (see the preceding description of fun), or else fminunc uses the medium-scale algorithm. For information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.
MaxFunEvals	Maximum number of function evaluations allowed, a positive integer. The default value is 100*numberOfVariables.
MaxIter	Maximum number of iterations allowed, a positive integer. The default value is 400.
OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none ([]). See “Output Function” on page 9-21.

PlotFcns	<p>Plots various measures of progress while the algorithm executes. Select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]).</p> <ul style="list-style-type: none">• @optimplotx plots the current point.• @optimplotfunccount plots the function count.• @optimplotfval plots the function value.• @optimplotstepsize plots the step size.• @optimplotfirstorderopt plots the first-order optimality measure. <p>For information on writing a custom plot function, see “Plot Functions” on page 9-30.</p>
TolFun	<p>Termination tolerance on the function value, a positive scalar. The default is $1e-6$.</p>
TolX	<p>Termination tolerance on x, a positive scalar. The default value is $1e-6$.</p>
TypicalX	<p>Typical x values. The number of elements in TypicalX is equal to the number of elements in x0, the starting point. The default value is <code>ones(numberofvariables,1)</code>. <code>fminunc</code> uses TypicalX for scaling finite differences for gradient estimation.</p> <p>The trust-region algorithm uses TypicalX only for the DerivativeCheck option.</p>

trust-region Algorithm Only

The trust-region algorithm uses the following options:

Hessian If 'on', fminunc uses a user-defined Hessian for the objective function. The Hessian is either defined in the objective function (see fun), or is indirectly defined when using HessMult.

If 'off' (default), fminunc approximates the Hessian using finite differences.

HessMult Function handle for Hessian multiply function. For large-scale structured problems, this function computes the Hessian matrix product $H*Y$ without actually forming H . The function is of the form

$W = \text{hmfun}(\text{Hinfo}, Y)$

where **Hinfo** contains the matrix used to compute $H*Y$.

The first argument must be the same as the third argument returned by the objective function fun, for example by

$[f, g, \text{Hinfo}] = \text{fun}(x)$

Y is a matrix that has the same number of rows as there are dimensions in the problem. $W = H*Y$ although H is not formed explicitly. fminunc uses **Hinfo** to compute the preconditioner. See “Passing Extra Parameters” on page 2-53 for information on how to supply values for any additional parameters hmfun needs.

Note 'Hessian' must be set to 'on' for fminunc to pass **Hinfo** from fun to hmfun.

See “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80 for an example.

HessPattern

Sparsity pattern of the Hessian for finite differencing. Set `HessPattern(i,j) = 1` when you can have $\partial^2 \text{fun} / \partial x(i) \partial x(j) \neq 0$. Otherwise, set `HessPattern(i,j) = 0`.

Use `HessPattern` when it is inconvenient to compute the Hessian matrix H in `fun`, but you can determine (say, by inspection) when the i th component of the gradient of `fun` depends on $x(j)$. `fminunc` can approximate H via sparse finite differences (of the gradient) if you provide the *sparsity structure* of H — i.e., locations of the nonzeros — as the value for `HessPattern`.

In the worst case, when the structure is unknown, do not set `HessPattern`. The default behavior is as if `HessPattern` is a dense matrix of ones. Then `fminunc` computes a full finite-difference approximation in each iteration. This can be very expensive for large problems, so it is usually better to determine the sparsity structure.

MaxPCGIter

Maximum number of PCG (preconditioned conjugate gradient) iterations, a positive scalar. The default is `max(1, floor(numberOfVariables/2))`. For more information, see “Algorithms” on page 10-109.

<code>PrecondBandWidth</code>	Upper bandwidth of preconditioner for PCG, a nonnegative integer. By default, <code>fminunc</code> uses diagonal preconditioning (upper bandwidth of 0). For some problems, increasing the bandwidth reduces the number of PCG iterations. Setting <code>PrecondBandWidth</code> to <code>Inf</code> uses a direct factorization (Cholesky) rather than the conjugate gradients (CG). The direct factorization is computationally more expensive than CG, but produces a better quality step towards the solution.
<code>TolPCG</code>	Termination tolerance on the PCG iteration, a positive scalar. The default is 0.1.

quasi-newton Algorithm Only

The quasi-newton algorithm uses the following options:

<code>HessUpdate</code>	Method for choosing the search direction in the Quasi-Newton algorithm. The choices are: <ul style="list-style-type: none">• 'bfgs', the default• 'dfp'• 'steepdesc' See "Hessian Update" on page 6-13 for a description of these methods.
<code>InitialHessMatrix</code>	Initial quasi-Newton matrix. This option is only available if you set <code>InitialHessType</code> to 'user-supplied'. In that case, you can set <code>InitialHessMatrix</code> to one of the following: <ul style="list-style-type: none">• A positive scalar — The initial matrix is the scalar times the identity.

	<ul style="list-style-type: none"> • A vector of positive values— The initial matrix is a diagonal matrix with the entries of the vector on the diagonal. This vector should be the same size as the <code>x0</code> vector, the initial point.
<code>InitialHesType</code>	Initial quasi-Newton matrix type. The options are: <ul style="list-style-type: none"> • <code>'identity'</code> • <code>'scaled-identity'</code>, the default • <code>'user-supplied'</code> — See <code>InitialHessMatrix</code>
<code>ObjectiveLimit</code>	A tolerance (stopping criterion) that is a scalar. If the objective function value at an iteration is less than or equal to <code>ObjectiveLimit</code> , the iterations halt, since the problem is presumably unbounded. The default value is <code>-1e20</code> .

Examples

Minimize the function $f(x) = 3x_1^2 + 2x_1x_2 + x_2^2$.

Create a file `myfun.m`:

```
function f = myfun(x)
f = 3*x(1)^2 + 2*x(1)*x(2) + x(2)^2;    % Cost function
```

Then call `fminunc` to find a minimum of `myfun` near `[1,1]`:

```
x0 = [1,1];
[x,fval] = fminunc(@myfun,x0);
```

After a few iterations, `fminunc` returns the solution, `x`, and the value of the function at `x`, `fval`:

```
x,fval
```

fminunc

```
x =  
 1.0e-006 *  
 0.2541   -0.2029
```

```
fval =  
 1.3173e-013
```

To minimize this function with the gradient provided, modify `myfun.m` so the gradient is the second output argument:

```
function [f,g] = myfun(x)  
f = 3*x(1)^2 + 2*x(1)*x(2) + x(2)^2;    % Cost function  
if nargin > 1  
    g(1) = 6*x(1)+2*x(2);  
    g(2) = 2*x(1)+2*x(2);  
end
```

Indicate that the gradient value is available by creating optimization options with the `GradObj` option set to 'on' using `optimoptions`.

```
options = optimoptions('fminunc','GradObj','on');  
x0 = [1,1];  
[x,fval] = fminunc(@myfun,x0,options);
```

After several iterations `fminunc` returns the solution, `x`, and the value of the function at `x`, `fval`:

```
x,fval
```

```
x =  
 1.0e-015 *  
 0.1110   -0.8882
```

```
fval =  
 6.2862e-031
```

To minimize the function $f(x) = \sin(x) + 3$ using an anonymous function

```
f = @(x)sin(x)+3;
x = fminunc(f,4);
```

fminunc returns a solution

x

```
x =
    4.7124
```

Notes

fminunc is not the preferred choice for solving problems that are sums of squares, that is, of the form

$$\min_x \|f(x)\|_2^2 = \min_x (f_1(x)^2 + f_2(x)^2 + \dots + f_n(x)^2)$$

Instead use the lsqnonlin function, which has been optimized for problems of this form.

To use the trust-region method, you must provide the gradient in fun (and set the GradObj option to 'on' using optimoptions). A warning is given if no gradient is provided and the Algorithm option is 'trust-region'.

Algorithms

Trust Region Algorithm

By default fminunc chooses the trust-region algorithm if you supply the gradient in fun and set GradObj to 'on' using optimoptions. This algorithm is a subspace trust-region method and is based on the interior-reflective Newton method described in [2] and [3]. Each iteration involves the approximate solution of a large linear system using the method of preconditioned conjugate gradients (PCG). See “fminunc trust-region Algorithm” on page 6-5, “Trust-Region Methods for Nonlinear Minimization” on page 6-5 and “Preconditioned Conjugate Gradient Method” on page 6-8pt.

Quasi-Newton Algorithm

The quasi-newton algorithm uses the BFGS Quasi-Newton method with a cubic line search procedure. This quasi-Newton method uses the BFGS ([1],[5],[8], and [9]) formula for updating the approximation of the Hessian matrix. You can select the DFP ([4],[6], and [7]) formula, which approximates the inverse Hessian matrix, by setting the `HessUpdate` option to 'dfp' (and the `Algorithm` option to 'quasi-newton'). You can select a steepest descent method by setting `HessUpdate` to 'steepdesc' (and `Algorithm` to 'quasi-newton'), although this is not recommended.

Limitations

The function to be minimized must be continuous. `fminunc` might only give local solutions.

`fminunc` only minimizes over the real numbers, that is, x must only consist of real numbers and $f(x)$ must only return real numbers. When x has complex variables, they must be split into real and imaginary parts.

Trust Region Optimization

To use the trust-region algorithm, you must supply the gradient in `fun` (and `GradObj` must be set 'on' in options).

Trust Region Algorithm Coverage and Requirements

Additional Information Needed	For Large Problems
Must provide gradient for $f(x)$ in <code>fun</code> .	<ul style="list-style-type: none">• Provide sparsity structure of the Hessian, or compute the Hessian in <code>fun</code>.• The Hessian should be sparse.

References

- [1] Broyden, C.G., “The Convergence of a Class of Double-Rank Minimization Algorithms,” *Journal Inst. Math. Applic.*, Vol. 6, pp. 76-90, 1970.
- [2] Coleman, T.F. and Y. Li, “An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds,” *SIAM Journal on Optimization*, Vol. 6, pp. 418-445, 1996.
- [3] Coleman, T.F. and Y. Li, “On the Convergence of Reflective Newton Methods for Large-Scale Nonlinear Minimization Subject to Bounds,” *Mathematical Programming*, Vol. 67, Number 2, pp. 189-224, 1994.
- [4] Davidon, W.C., “Variable Metric Method for Minimization,” *A.E.C. Research and Development Report*, ANL-5990, 1959.
- [5] Fletcher, R., “A New Approach to Variable Metric Algorithms,” *Computer Journal*, Vol. 13, pp. 317-322, 1970.
- [6] Fletcher, R., “Practical Methods of Optimization,” Vol. 1, *Unconstrained Optimization*, John Wiley and Sons, 1980.
- [7] Fletcher, R. and M.J.D. Powell, “A Rapidly Convergent Descent Method for Minimization,” *Computer Journal*, Vol. 6, pp. 163-168, 1963.
- [8] Goldfarb, D., “A Family of Variable Metric Updates Derived by Variational Means,” *Mathematics of Computing*, Vol. 24, pp. 23-26, 1970.
- [9] Shanno, D.F., “Conditioning of Quasi-Newton Methods for Function Minimization,” *Mathematics of Computing*, Vol. 24, pp. 647-656, 1970.

See Also

fminsearch | optimoptions | optimtool

How To

- `function_handle`
- “Anonymous Functions”
- “Unconstrained Optimization”

fseminf

Purpose Find minimum of semi-infinitely constrained multivariable nonlinear function

Equation Finds the minimum of a problem specified by

$$\min_x f(x) \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ lb \leq x \leq ub, \\ c(x) \leq 0, \\ ceq(x) = 0, \\ K_i(x, w_i) \leq 0, \quad 1 \leq i \leq n. \end{cases}$$

b and beq are vectors, A and Aeq are matrices, $c(x)$, $ceq(x)$, and $K_i(x, w_i)$ are functions that return vectors, and $f(x)$ is a function that returns a scalar. $f(x)$, $c(x)$, and $ceq(x)$ can be nonlinear functions. The vectors (or matrices) $K_i(x, w_i) \leq 0$ are continuous functions of both x and an additional set of variables w_1, w_2, \dots, w_n . The variables w_1, w_2, \dots, w_n are vectors of, at most, length two.

x , lb , and ub can be passed as vectors or matrices; see “Matrix Arguments” on page 2-32.

Syntax

```
x = fseminf(fun,x0,ntheta,seminfcon)
x = fseminf(fun,x0,ntheta,seminfcon,A,b)
x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq)
x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq,lb,ub)
x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq,lb,ub,options)
x = fseminf(problem)
[x,fval] = fseminf(...)
[x,fval,exitflag] = fseminf(...)
[x,fval,exitflag,output] = fseminf(...)
[x,fval,exitflag,output,lambda] = fseminf(...)
```

Description `fseminf` finds a minimum of a semi-infinitely constrained scalar function of several variables, starting at an initial estimate. The aim is to minimize $f(x)$ so the constraints hold for all possible values of

$w_i \mathcal{N}^1$ (or $w_i \mathcal{N}^2$). Because it is impossible to calculate all possible values of $K_i(x, w_i)$, a region must be chosen for w_i over which to calculate an appropriately sampled set of values.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective function and nonlinear constraint functions, if necessary.

`x = fseminf(fun,x0,ntheta,seminfcon)` starts at `x0` and finds a minimum of the function `fun` constrained by `ntheta` semi-infinite constraints defined in `seminfcon`.

`x = fseminf(fun,x0,ntheta,seminfcon,A,b)` also tries to satisfy the linear inequalities $A \cdot x \leq b$.

`x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq)` minimizes subject to the linear equalities $Aeq \cdot x = beq$ as well. Set `A = []` and `b = []` if no inequalities exist.

`x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq,lb,ub)` defines a set of lower and upper bounds on the design variables in `x`, so that the solution is always in the range $lb \leq x \leq ub$.

Note See “Iterations Can Violate Constraints” on page 2-34.

`x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq,lb,ub,options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options.

`x = fseminf(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-114.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

fseminf

`[x,fval] = fseminf(...)` returns the value of the objective function `fun` at the solution `x`.

`[x,fval,exitflag] = fseminf(...)` returns a value `exitflag` that describes the exit condition.

`[x,fval,exitflag,output] = fseminf(...)` returns a structure `output` that contains information about the optimization.

`[x,fval,exitflag,output,lambda] = fseminf(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the output `fval` is `[]`.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `fseminf`. This section provides function-specific details for `fun`, `ntheta`, `options`, `seminfcon`, and `problem`:

`fun`

The function to be minimized. `fun` is a function that accepts a vector `x` and returns a scalar `f`, the objective function evaluated at `x`. The function `fun` can be specified as a function handle for a file

```
x = fseminf(@myfun,x0,ntheta,seminfcon)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...           % Compute function value at x
```

`fun` can also be a function handle for an anonymous function.

```
fun = @(x)sin(x'*x);
```

If the gradient of `fun` can also be computed *and* the `GradObj` option is `'on'`, as set by


```
options = optimoptions('fseminf','GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `g`, a vector, at `x`.

`ntheta` The number of semi-infinite constraints.

`options` “Options” on page 10-118 provides the function-specific details for the options values.

`seminfcon` The function that computes the vector of nonlinear inequality constraints, `c`, a vector of nonlinear equality constraints, `ceq`, and `ntheta` semi-infinite constraints (vectors or matrices) `K1`, `K2`, ..., `Kntheta` evaluated over an interval `S` at the point `x`. The function `seminfcon` can be specified as a function handle.

```
x = fseminf(@myfun,x0,ntheta,@myinfcon)
```

where `myinfcon` is a MATLAB function such as

```
function [c,ceq,K1,K2,...,Kntheta,S] = myinfcon(x,S)
% Initial sampling interval
if isnan(S(1,1)),
    S = ...% S has ntheta rows and 2 columns
end
w1 = ...% Compute sample set
w2 = ...% Compute sample set
...
wntheta = ... % Compute sample set
K1 = ... % 1st semi-infinite constraint at x and w
K2 = ... % 2nd semi-infinite constraint at x and w
...
Kntheta = ...% Last semi-infinite constraint at x and w
c = ...      % Compute nonlinear inequalities at x
ceq = ...    % Compute the nonlinear equalities at x
```

`S` is a recommended sampling interval, which might or might not be used. Return `[]` for `c` and `ceq` if no such constraints exist.

The vectors or matrices $K_1, K_2, \dots, K_{n\theta}$ contain the semi-infinite constraints evaluated for a sampled set of values for the independent variables $w_1, w_2, \dots, w_{n\theta}$, respectively. The two-column matrix, S , contains a recommended sampling interval for values of $w_1, w_2, \dots, w_{n\theta}$, which are used to evaluate $K_1, K_2, \dots, K_{n\theta}$. The i th row of S contains the recommended sampling interval for evaluating K_i . When K_i is a vector, use only $S(i, 1)$ (the second column can be all zeros). When K_i is a matrix, $S(i, 2)$ is used for the sampling of the rows in K_i , $S(i, 1)$ is used for the sampling interval of the columns of K_i (see “Two-Dimensional Semi-Infinite Constraint” on page 6-103). On the first iteration S is NaN, so that some initial sampling interval must be determined by `seminfcon`.

Note Because Optimization Toolbox functions only accept inputs of type `double`, user-supplied objective and nonlinear constraint functions must return outputs of type `double`.

“Passing Extra Parameters” on page 2-53 explains how to parameterize `seminfcon`, if necessary. “Example of Creating Sampling Points” on page 6-49 contains an example of both one- and two-dimensional sampling points.

problem	objective	Objective function
	x0	Initial point for x
	ntheta	Number of semi-infinite constraints
	seminfcon	Semi-infinite constraint function
	Aineq	Matrix for linear inequality constraints
	bineq	Vector for linear inequality constraints
	Aeq	Matrix for linear equality constraints
	beq	Vector for linear equality constraints
	lb	Vector of lower bounds

ub	Vector of upper bounds
solver	'fseminf'
options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `fseminf`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

`exitflag` Integer identifying the reason the algorithm terminated. The following lists the values of `exitflag` and the corresponding reasons the algorithm terminated.

1	Function converged to a solution <code>x</code> .
4	Magnitude of the search direction was less than the specified tolerance and constraint violation was less than <code>options.TolCon</code> .
5	Magnitude of directional derivative was less than the specified tolerance and constraint violation was less than <code>options.TolCon</code> .
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .
-1	Algorithm was terminated by the output function.
-2	No feasible point was found.

fseminf

<code>lambda</code>	Structure containing the Lagrange multipliers at the solution <code>x</code> (separated by constraint type). The fields of the structure are
<code>lower</code>	Lower bounds <code>lb</code>
<code>upper</code>	Upper bounds <code>ub</code>
<code>ineqlin</code>	Linear inequalities
<code>eqlin</code>	Linear equalities
<code>ineqnonlin</code>	Nonlinear inequalities
<code>eqnonlin</code>	Nonlinear equalities
<code>output</code>	Structure containing information about the optimization. The fields of the structure are
<code>iterations</code>	Number of iterations taken
<code>funcCount</code>	Number of function evaluations
<code>lssteplength</code>	Size of line search step relative to search direction
<code>stepsize</code>	Final displacement in <code>x</code>
<code>algorithm</code>	Optimization algorithm used
<code>constrviolation</code>	Maximum of constraint functions
<code>firstorderopt</code>	Measure of first-order optimality
<code>message</code>	Exit message

Options

Optimization options used by `fseminf`. Use `optimoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information.

DerivativeCheck	Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives. The choices are 'on' or the default 'off'.
Diagnostics	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default 'off'.
DiffMaxChange	Maximum change in variables for finite-difference gradients (a positive scalar). The default is Inf.
DiffMinChange	Minimum change in variables for finite-difference gradients (a positive scalar). The default is 0.
Display	Level of display: <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'iter' displays output at each iteration, and gives the default exit message.• 'iter-detailed' displays output at each iteration, and gives the technical exit message.• 'notify' displays output only if the function does not converge, and gives the default exit message.• 'notify-detailed' displays output only if the function does not converge, and gives the technical exit message.• 'final' (default) displays just the final output, and gives the default exit message.

- 'final-detailed' displays just the final output, and gives the technical exit message.

FinDiffRelStep

Scalar or vector step size factor. When you set `FinDiffRelStep` to a vector `v`, forward finite differences `delta` are

```
delta =  
v.*sign(x).*max(abs(x),TypicalX);
```

and central finite differences are

```
delta = v.*max(abs(x),TypicalX);
```

Scalar `FinDiffRelStep` expands to a vector. The default is `sqrt(eps)` for forward finite differences, and `eps^(1/3)` for central finite differences.

FinDiffType

Finite differences, used to estimate gradients, are either 'forward' (the default), or 'central' (centered). 'central' takes twice as many function evaluations, but should be more accurate.

The algorithm is careful to obey bounds when estimating both types of finite differences. So, for example, it could take a backward, rather than a forward, difference to avoid evaluating at a point outside bounds.

FunValCheck	Check whether objective function and constraints values are valid. 'on' displays an error when the objective function or constraints return a value that is complex, Inf, or NaN. The default 'off' displays no error.
GradObj	Gradient for the objective function defined by the user. See the preceding description of fun to see how to define the gradient in fun. Set to 'on' to have fseminf use a user-defined gradient of the objective function. The default 'off' causes fseminf to estimate gradients using finite differences.
MaxFunEvals	Maximum number of function evaluations allowed, a positive integer. The default is 100*numberOfVariables.
MaxIter	Maximum number of iterations allowed, a positive integer. The default is 400.
MaxSQPIter	Maximum number of SQP iterations allowed, a positive integer. The default is 10*max(numberOfVariables, numberOfInequalities + numberOfBounds).
OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none ([]). See “Output Function” on page 9-21.

PlotFcns	<p>Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]):</p> <ul style="list-style-type: none">• @optimplotx plots the current point.• @optimplotfunccount plots the function count.• @optimplotfval plots the function value.• @optimplotconstrviolation plots the maximum constraint violation.• @optimplotstepsize plots the step size.• @optimplotfirstorderopt plots the first-order optimality measure. <p>For information on writing a custom plot function, see “Plot Functions” on page 9-30.</p>
RelLineSrchBnd	<p>Relative bound (a real nonnegative scalar value) on the line search step length such that the total displacement in x satisfies $\Delta x(i) \leq \text{relLineSrchBnd} \cdot \max(x(i) , typicalx(i))$. This option provides control over the magnitude of the displacements in x for cases in which the solver takes steps that <code>fseminf</code> considers too large. The default is no bounds ([]).</p>
RelLineSrchBndDuration	<p>Number of iterations for which the bound specified in <code>RelLineSrchBnd</code> should be active (default is 1)</p>

TolCon	Termination tolerance on the constraint violation, a positive scalar. The default is 1e-6.
TolConSQP	Termination tolerance on inner iteration SQP constraint violation, a positive scalar. The default is 1e-6.
TolFun	Termination tolerance on the function value, a positive scalar. The default is 1e-4.
TolX	Termination tolerance on x, a positive scalar. The default value is 1e-4.
TypicalX	Typical x values. The number of elements in TypicalX is equal to the number of elements in x0, the starting point. The default value is ones(numberofvariables,1). fseminf uses TypicalX for scaling finite differences for gradient estimation.

Notes

The optimization routine `fseminf` might vary the recommended sampling interval, `S`, set in `seminfcon`, during the computation because values other than the recommended interval might be more appropriate for efficiency or robustness. Also, the finite region w_i , over which $K_i(x, w_i)$ is calculated, is allowed to vary during the optimization, provided that it does not result in significant changes in the number of local minima in $K_i(x, w_i)$.

Examples

This example minimizes the function

$$(x - 1)^2,$$

subject to the constraints

$$0 \leq x \leq 2$$
$$g(x, t) = (x - 1/2) - (t - 1/2)^2 \leq 0 \text{ for all } 0 \leq t \leq 1.$$

The unconstrained objective function is minimized at $x = 1$. However, the constraint,

$$g(x, t) \leq 0 \text{ for all } 0 \leq t \leq 1,$$

implies $x \leq 1/2$. You can see this by noticing that $(t - 1/2)^2 \geq 0$, so

$$\max_t g(x, t) = (x - 1/2).$$

Therefore

$$\max_t g(x, t) \leq 0 \text{ when } x \leq 1/2.$$

To solve this problem using `fseminf`:

- 1 Write the objective function as an anonymous function:

```
objfun = @(x) (x-1)^2;
```

- 2 Write the semi-infinite constraint function, which includes the nonlinear constraints (`[]` in this case), initial sampling interval for t (0 to 1 in steps of 0.01 in this case), and the semi-infinite constraint function $g(x, t)$:

```
function [c, ceq, K1, s] = seminfcon(x,s)

% No finite nonlinear inequality and equality constraints
c = [];
ceq = [];

% Sample set
if isnan(s)
    % Initial sampling interval
    s = [0.01 0];
end
```

```
t = 0:s(1):1;

% Evaluate the semi-infinite constraint
K1 = (x - 0.5) - (t - 0.5).^2;
```

3 Call `fseminf` with initial point 0.2, and view the result:

```
x = fseminf(objfun,0.2,1,@seminfcon)
```

Local minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.

```
Active inequalities (to within options.TolCon = 1e-006):
   lower      upper   ineqlin   ineqnonlin
           1
```

```
x =
    0.5000
```

Algorithms

`fseminf` uses cubic and quadratic interpolation techniques to estimate peak values in the semi-infinite constraints. The peak values are used to form a set of constraints that are supplied to an SQP method as in the `fmincon` function. When the number of constraints changes, Lagrange multipliers are reallocated to the new set of constraints.

The recommended sampling interval calculation uses the difference between the interpolated peak values and peak values appearing in the data set to estimate whether the function needs to take more or fewer points. The function also evaluates the effectiveness of the interpolation by extrapolating the curve and comparing it to other points in the curve. The recommended sampling interval is decreased when the peak values are close to constraint boundaries, i.e., zero.

fseminf

For more details on the algorithm used and the types of procedures displayed under the Procedures heading when the Display option is set to 'iter' with `optimoptions`, see also “SQP Implementation” on page 6-35. For more details on the `fseminf` algorithm, see “fseminf Problem Formulation and Algorithm” on page 6-47.

Limitations

The function to be minimized, the constraints, and semi-infinite constraints, must be continuous functions of x and w . `fseminf` might only give local solutions.

When the problem is not feasible, `fseminf` attempts to minimize the maximum constraint value.

See Also

`fmincon` | `optimoptions` | `optimtool`

How To

- `function_handle`
- “fseminf Problem Formulation and Algorithm” on page 6-47
- “Constrained Optimization”

Purpose Solve system of nonlinear equations

Equation Solves a problem specified by

$$F(x) = 0$$

for x , where $F(x)$ is a function that returns a vector value.

x is a vector or a matrix; see “Matrix Arguments” on page 2-32.

Syntax

```
x = fsolve(fun,x0)
x = fsolve(fun,x0,options)
x = fsolve(problem)
[x,fval] = fsolve(fun,x0)
[x,fval,exitflag] = fsolve(...)
[x,fval,exitflag,output] = fsolve(...)
[x,fval,exitflag,output,jacobian] = fsolve(...)
```

Description

fsolve finds a root (zero) of a system of nonlinear equations.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the system of equations, if necessary.

`x = fsolve(fun,x0)` starts at `x0` and tries to solve the equations described in `fun`.

`x = fsolve(fun,x0,options)` solves the equations with the optimization options specified in `options`. Use `optimoptions` to set these options.

`x = fsolve(problem)` solves `problem`, where `problem` is a structure described in “Input Arguments” on page 10-128.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = fsolve(fun,x0)` returns the value of the objective function `fun` at the solution `x`.

`[x,fval,exitflag] = fsolve(...)` returns a value `exitflag` that describes the exit condition.

`[x,fval,exitflag,output] = fsolve(...)` returns a structure `output` that contains information about the optimization.

`[x,fval,exitflag,output,jacobian] = fsolve(...)` returns the Jacobian of `fun` at the solution `x`.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `fsolve`. This section provides function-specific details for `fun` and problem:

fun The nonlinear system of equations to solve. `fun` is a function that accepts a vector `x` and returns a vector `F`, the nonlinear equations evaluated at `x`. The function `fun` can be specified as a function handle for a file

```
x = fsolve(@myfun,x0)
```

where `myfun` is a MATLAB function such as

```
function F = myfun(x)
F = ...           % Compute function values at x
```

`fun` can also be a function handle for an anonymous function.

```
x = fsolve(@(x)sin(x.*x),x0);
```

If the user-defined values for `x` and `F` are matrices, they are converted to a vector using linear indexing.

If the Jacobian can also be computed *and* the Jacobian option is 'on', set by

```
options = optimoptions('fsolve','Jacobian','on')
```

the function `fun` must return, in a second output argument, the Jacobian value `J`, a matrix, at `x`.

If `fun` returns a vector (matrix) of `m` components and `x` has length `n`, where `n` is the length of `x0`, the Jacobian `J` is an `m`-by-`n` matrix where `J(i,j)` is the partial derivative of `F(i)` with respect to `x(j)`. (The Jacobian `J` is the transpose of the gradient of `F`.)

<code>problem</code>	<code>objective</code>	Objective function
	<code>x0</code>	Initial point for <code>x</code>
	<code>solver</code>	' <code>fsolve</code> '
	<code>options</code>	Options created with <code>optimoptions</code>

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `fsolve`. For more information on the output headings for `fsolve`, see “Function-Specific Headings” on page 3-18.

This section provides function-specific details for `exitflag` and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.
1	Function converged to a solution <code>x</code> .
2	Change in <code>x</code> was smaller than the specified tolerance.
3	Change in the residual was smaller than the specified tolerance.
4	Magnitude of search direction was smaller than the specified tolerance.
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .

	-1	Output function terminated the algorithm.
	-2	Algorithm appears to be converging to a point that is not a root.
	-3	Trust region radius became too small (trust-region-dogleg algorithm).
	-4	Line search cannot sufficiently decrease the residual along the current search direction.
output		Structure containing information about the optimization. The fields of the structure are
	iterations	Number of iterations taken
	funcCount	Number of function evaluations
	algorithm	Optimization algorithm used
	cgiterations	Total number of PCG iterations (trust-region-reflective algorithm only)
	stepsize	Final displacement in x (Levenberg-Marquardt algorithm)
	firstorderopt	Measure of first-order optimality (dogleg or trust-region-reflective algorithm, [] for others)
	message	Exit message

Options

Optimization options used by `fsolve`. Some options apply to all algorithms, some are only relevant when using the trust-region-reflective algorithm, and others are only relevant when using the other algorithms. Use `optimoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information.

All Algorithms

All algorithms use the following options:

Algorithm	<p>Choose between 'trust-region-dogleg' (default), 'trust-region-reflective', and 'levenberg-marquardt'. Set the initial Levenberg-Marquardt parameter λ by setting Algorithm to a cell array such as {'levenberg-marquardt', .005}. The default $\lambda = 0.01$.</p> <p>The Algorithm option specifies a preference for which algorithm to use. It is only a preference because for the trust-region-reflective algorithm, the nonlinear system of equations cannot be underdetermined; that is, the number of equations (the number of elements of F returned by fun) must be at least as many as the length of x. Similarly, for the trust-region-dogleg algorithm, the number of equations must be the same as the length of x. fsolve uses the Levenberg-Marquardt algorithm when the selected algorithm is unavailable. For more information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.</p>
DerivativeCheck	Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives. The choices are 'on' or the default 'off'.
Diagnostics	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default 'off'.
DiffMaxChange	Maximum change in variables for finite-difference gradients (a positive scalar). The default is Inf.

DiffMinChange Minimum change in variables for finite-difference gradients (a positive scalar). The default is 0.

Display Level of display:

- 'off' or 'none' displays no output.
- 'iter' displays output at each iteration, and gives the default exit message.
- 'iter-detailed' displays output at each iteration, and gives the technical exit message.
- 'final' (default) displays just the final output, and gives the default exit message.
- 'final-detailed' displays just the final output, and gives the technical exit message.

FinDiffRelStep Scalar or vector step size factor. When you set **FinDiffRelStep** to a vector v , forward finite differences δ are

$$\delta = v \cdot \text{sign}(x) \cdot \max(\text{abs}(x), \text{TypicalX});$$

and central finite differences are

$$\delta = v \cdot \max(\text{abs}(x), \text{TypicalX});$$

Scalar **FinDiffRelStep** expands to a vector. The default is $\text{sqrt}(\text{eps})$ for forward finite differences, and $\text{eps}^{(1/3)}$ for central finite differences.

<code>FinDiffType</code>	<p>Finite differences, used to estimate gradients, are either 'forward' (default), or 'central' (centered). 'central' takes twice as many function evaluations, but should be more accurate.</p> <p>The algorithm is careful to obey bounds when estimating both types of finite differences. So, for example, it could take a backward, rather than a forward, difference to avoid evaluating at a point outside bounds.</p>
<code>FunValCheck</code>	<p>Check whether objective function values are valid. 'on' displays an error when the objective function returns a value that is complex, Inf, or NaN. The default, 'off', displays no error.</p>
<code>Jacobian</code>	<p>If 'on', <code>fsolve</code> uses a user-defined Jacobian (defined in <code>fun</code>), or Jacobian information (when using <code>JacobMult</code>), for the objective function. If 'off' (default), <code>fsolve</code> approximates the Jacobian using finite differences.</p>
<code>MaxFunEvals</code>	<p>Maximum number of function evaluations allowed, a positive integer. The default is <code>100*numberOfVariables</code>.</p>
<code>MaxIter</code>	<p>Maximum number of iterations allowed, a positive integer. The default is 400.</p>
<code>OutputFcn</code>	<p>Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none (<code>[]</code>). See “Output Function” on page 9-21.</p>

PlotFcns	<p>Plots various measures of progress while the algorithm executes. Select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none (<code>[]</code>):</p> <ul style="list-style-type: none">• <code>@optimplotx</code> plots the current point.• <code>@optimplotfunccount</code> plots the function count.• <code>@optimplotfval</code> plots the function value.• <code>@optimplotresnorm</code> plots the norm of the residuals.• <code>@optimplotstepsize</code> plots the step size.• <code>@optimplotfirstorderopt</code> plots the first-order optimality measure. <p>For information on writing a custom plot function, see “Plot Functions” on page 9-30.</p>
TolFun	<p>Termination tolerance on the function value, a positive scalar. The default is <code>1e-6</code>.</p>
TolX	<p>Termination tolerance on <code>x</code>, a positive scalar. The default is <code>1e-6</code>.</p>
TypicalX	<p>Typical <code>x</code> values. The number of elements in <code>TypicalX</code> is equal to the number of elements in <code>x0</code>, the starting point. The default value is <code>ones(numberofvariables,1)</code>. <code>fsolve</code> uses <code>TypicalX</code> for scaling finite differences for gradient estimation.</p> <p>The trust-region-dogleg algorithm uses <code>TypicalX</code> as the diagonal terms of a scaling matrix.</p>

Trust-Region-Reflective Algorithm Only

The trust-region-reflective algorithm uses the following options:

JacobMult

Function handle for Jacobian multiply function. For large-scale structured problems, this function computes the Jacobian matrix product $J*Y$, $J'*Y$, or $J'*(J*Y)$ without actually forming J . The function is of the form

$$W = \text{jmfun}(J\text{info}, Y, \text{flag})$$

where $J\text{info}$ contains a matrix used to compute $J*Y$ (or $J'*Y$, or $J'*(J*Y)$). The first argument $J\text{info}$ must be the same as the second argument returned by the objective function fun , for example, in

$$[F, J\text{info}] = \text{fun}(x)$$

Y is a matrix that has the same number of rows as there are dimensions in the problem. flag determines which product to compute:

- If $\text{flag} == 0$, $W = J'*(J*Y)$.
- If $\text{flag} > 0$, $W = J*Y$.
- If $\text{flag} < 0$, $W = J'*Y$.

In each case, J is not formed explicitly. `fsolve` uses $J\text{info}$ to compute the preconditioner. See “Passing Extra Parameters” on page 2-53 for information on how to supply values for any additional parameters `jmfun` needs.

Note 'Jacobian' must be set to 'on' for `fsolve` to pass $J\text{info}$ from `fun` to `jmfun`.

See “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80 for a similar example.

JacobPattern

Sparsity pattern of the Jacobian for finite differencing. Set `JacobPattern(i, j) = 1` when `fun(i)` depends on `x(j)`. Otherwise, set `JacobPattern(i, j) = 0`. In other words, `JacobPattern(i, j) = 1` when you can have $\partial \text{fun}(i) / \partial x(j) \neq 0$.

Use `JacobPattern` when it is inconvenient to compute the Jacobian matrix `J` in `fun`, though you can determine (say, by inspection) when `fun(i)` depends on `x(j)`. `fsolve` can approximate `J` via sparse finite differences when you give `JacobPattern`.

In the worst case, if the structure is unknown, do not set `JacobPattern`. The default behavior is as if `JacobPattern` is a dense matrix of ones. Then `fsolve` computes a full finite-difference approximation in each iteration. This can be very expensive for large problems, so it is usually better to determine the sparsity structure.

MaxPCGIter

Maximum number of PCG (preconditioned conjugate gradient) iterations, a positive scalar. The default is `max(1, floor(numberOfVariables/2))`. For more information, see “Algorithms” on page 10-140.

PrecondBandWidth	Upper bandwidth of preconditioner for PCG, a nonnegative integer. The default PrecondBandWidth is Inf, which means a direct factorization (Cholesky) is used rather than the conjugate gradients (CG). The direct factorization is computationally more expensive than CG, but produces a better quality step towards the solution. Set PrecondBandWidth to 0 for diagonal preconditioning (upper bandwidth of 0). For some problems, an intermediate bandwidth reduces the number of PCG iterations.
TolPCG	Termination tolerance on the PCG iteration, a positive scalar. The default is 0.1.

Levenberg-Marquardt Algorithm Only

The Levenberg-Marquardt algorithm uses the following option:

ScaleProblem	'Jacobian' can sometimes improve the convergence of a poorly scaled problem. The default is 'none'.
--------------	---

Examples

Example 1

This example solves the system of two equations and two unknowns:

$$\begin{aligned} 2x_1 - x_2 &= e^{-x_1} \\ -x_1 + 2x_2 &= e^{-x_2}. \end{aligned}$$

Rewrite the equations in the form $F(x) = 0$:

$$\begin{aligned} 2x_1 - x_2 - e^{-x_1} &= 0 \\ -x_1 + 2x_2 - e^{-x_2} &= 0. \end{aligned}$$

Start your search for a solution at $x_0 = [-5 \ -5]$.

First, write a file that computes F, the values of the equations at x.

```
function F = myfun(x)
F = [2*x(1) - x(2) - exp(-x(1));
     -x(1) + 2*x(2) - exp(-x(2))];
```

Save this function file as myfun.m somewhere on your MATLAB path.

Next, set up the initial point and options and call fsolve:

```
x0 = [-5; -5]; % Make a starting guess at the solution
options = optimoptions('fsolve','Display','iter'); % Option to display output
[x,fval] = fsolve(@myfun,x0,options) % Call solver
```

After several iterations, fsolve finds an answer:

Iteration	Func-count	f(x)	Norm of		
			step	First-order optimality	Trust-region radius
0	3	23535.6		2.29e+004	1
1	6	6001.72	1	5.75e+003	1
2	9	1573.51	1	1.47e+003	1
3	12	427.226	1	388	1
4	15	119.763	1	107	1
5	18	33.5206	1	30.8	1
6	21	8.35208	1	9.05	1
7	24	1.21394	1	2.26	1
8	27	0.016329	0.759511	0.206	2.5
9	30	3.51575e-006	0.111927	0.00294	2.5
10	33	1.64763e-013	0.00169132	6.36e-007	2.5

Equation solved.

fsolve completed because the vector of function values is near zero as measured by the default value of the function tolerance, and the problem appears regular as measured by the gradient.

x =


```

0.5671
0.5671

fval =
1.0e-006 *
-0.4059
-0.4059

```

Example 2

Find a matrix x that satisfies the equation

$$X * X * X = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix},$$

starting at the point $x = [1, 1; 1, 1]$.

First, write a file that computes the equations to be solved.

```

function F = myfun(x)
F = x*x*x-[1,2;3,4];

```

Save this function file as `myfun.m` somewhere on your MATLAB path.

Next, set up an initial point and options and call `fsolve`:

```

x0 = ones(2,2); % Make a starting guess at the solution
options = optimoptions('fsolve','Display','off'); % Turn off display
[x,Fval,exitflag] = fsolve(@myfun,x0,options)

```

The solution is

```

x =
-0.1291    0.8602
 1.2903    1.1612

Fval =
1.0e-009 *
-0.1621    0.0780
 0.1167   -0.0465

```

```
exitflag =  
    1
```

and the residual is close to zero.

```
sum(sum(Fval.*Fval))  
ans =  
    4.8133e-20
```

Notes

If the system of equations is linear, use `\` (matrix left division) for better speed and accuracy. For example, to find the solution to the following linear system of equations:

$$\begin{aligned}3x_1 + 11x_2 - 2x_3 &= 7 \\ x_1 + x_2 - 2x_3 &= 4 \\ x_1 - x_2 + x_3 &= 19.\end{aligned}$$

Formulate and solve the problem as

```
A = [ 3 11 -2; 1 1 -2; 1 -1 1];  
b = [ 7; 4; 19];  
x = A\b  
x =  
    13.2188  
    -2.3438  
     3.4375
```

Algorithms

The Levenberg-Marquardt and trust-region-reflective methods are based on the nonlinear least-squares algorithms also used in `lsqnonlin`. Use one of these methods if the system may not have a zero. The algorithm still returns a point where the residual is small. However, if the Jacobian of the system is singular, the algorithm might converge to a point that is not a solution of the system of equations (see “Limitations” on page 10-141 and “Diagnostics” on page 10-141 following).

- By default `fsolve` chooses the trust-region dogleg algorithm. The algorithm is a variant of the Powell dogleg method described in [8].

It is similar in nature to the algorithm implemented in [7]. See “Trust-Region Dogleg Method” on page 6-252.

- The trust-region-reflective algorithm is a subspace trust-region method and is based on the interior-reflective Newton method described in [1] and [2]. Each iteration involves the approximate solution of a large linear system using the method of preconditioned conjugate gradients (PCG). See “Trust-Region Reflective fsolve Algorithm” on page 6-249.
- The Levenberg-Marquardt method is described in references [4], [5], and [6]. See “Levenberg-Marquardt Method” on page 6-255.

Diagnostics

All Algorithms

fsolve may converge to a nonzero point and give this message:

```
Optimizer is stuck at a minimum that is not a root
Try again with a new starting guess
```

In this case, run fsolve again with other starting values.

Trust-Region-Dogleg Algorithm

For the trust-region dogleg method, fsolve stops if the step size becomes too small and it can make no more progress. fsolve gives this message:

```
The optimization algorithm can make no further progress:
Trust region radius less than 10*eps
```

In this case, run fsolve again with other starting values.

Limitations

The function to be solved must be continuous. When successful, fsolve only gives one root. fsolve may converge to a nonzero point, in which case, try other starting values.

fsolve only handles real variables. When x has complex variables, the variables must be split into real and imaginary parts.

Trust-Region-Reflective Algorithm

The preconditioner computation used in the preconditioned conjugate gradient part of the trust-region-reflective algorithm forms $J^T J$ (where J is the Jacobian matrix) before computing the preconditioner; therefore, a row of J with many nonzeros, which results in a nearly dense product $J^T J$, might lead to a costly solution process for large problems.

Trust-Region-Reflective Problem Coverage and Requirements

For Large Problems
<ul style="list-style-type: none">• Provide sparsity structure of the Jacobian or compute the Jacobian in fun.• The Jacobian should be sparse.

Number of Equations

The default trust-region dogleg method can only be used when the system of equations is square, i.e., the number of equations equals the number of unknowns. For the Levenberg-Marquardt method, the system of equations need not be square.

References

- [1] Coleman, T.F. and Y. Li, "An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds," *SIAM Journal on Optimization*, Vol. 6, pp. 418-445, 1996.
- [2] Coleman, T.F. and Y. Li, "On the Convergence of Reflective Newton Methods for Large-Scale Nonlinear Minimization Subject to Bounds," *Mathematical Programming*, Vol. 67, Number 2, pp. 189-224, 1994.
- [3] Dennis, J. E. Jr., "Nonlinear Least-Squares," *State of the Art in Numerical Analysis*, ed. D. Jacobs, Academic Press, pp. 269-312.
- [4] Levenberg, K., "A Method for the Solution of Certain Problems in Least-Squares," *Quarterly Applied Mathematics* 2, pp. 164-168, 1944.

- [5] Marquardt, D., “An Algorithm for Least-squares Estimation of Nonlinear Parameters,” *SIAM Journal Applied Mathematics*, Vol. 11, pp. 431-441, 1963.
- [6] Moré, J. J., “The Levenberg-Marquardt Algorithm: Implementation and Theory,” *Numerical Analysis*, ed. G. A. Watson, Lecture Notes in Mathematics 630, Springer Verlag, pp. 105-116, 1977.
- [7] Moré, J. J., B. S. Garbow, and K. E. Hillstom, *User Guide for MINPACK 1*, Argonne National Laboratory, Rept. ANL-80-74, 1980.
- [8] Powell, M. J. D., “A Fortran Subroutine for Solving Systems of Nonlinear Algebraic Equations,” *Numerical Methods for Nonlinear Algebraic Equations*, P. Rabinowitz, ed., Ch.7, 1970.

See Also

\ | lsqcurvefit | lsqnonlin | optimoptions | optimtool

How To

- `function_handle`
- “Anonymous Functions”
- “Systems of Nonlinear Equations”

fzero

Purpose

Root of nonlinear function

Syntax

```
x = fzero(fun,x0)
```

```
x = fzero(fun,x0,options)
```

```
x = fzero(problem)
```

```
[x,fval,exitflag,output] = fzero( __ )
```

Description

`x = fzero(fun,x0)` tries to find a point x where $\text{fun}(x) = 0$. This solution is where $\text{fun}(x)$ changes sign—`fzero` cannot find a root of a function such as x^2 .

`x = fzero(fun,x0,options)` uses options to modify the solution process.

`x = fzero(problem)` solves a root-finding problem specified by `problem`.

`[x,fval,exitflag,output] = fzero(__)` returns $\text{fun}(x)$ in the `fval` output, `exitflag` encoding the reason `fzero` stopped, and an output structure containing information on the solution process.

Input Arguments

fun - Function to solve

function handle

Function to solve, specified as a handle to a scalar-valued function. `fun` accepts a scalar x and returns a scalar $\text{fun}(x)$.

`fzero` solves $\text{fun}(x) = 0$. To solve an equation $\text{fun}(x) = c(x)$, instead solve $\text{fun2}(x) = \text{fun}(x) - c(x) = 0$.

To include extra parameters in your function, see the example “Root of Function with Extra Parameter” on page 10-150 and the section “Passing Extra Parameters” on page 2-53.

Example: @sin

Example: @myFunction

Example: @(x)(x-a)^5 - 3*x + a - 1

Data Types

function_handle

x0 - Initial value

scalar | 2-element vector

Initial value, specified as a real scalar or a 2-element real vector.

- Scalar — fzero begins at x_0 and tries to locate a point x_1 where $\text{fun}(x_1)$ has the opposite sign of $\text{fun}(x_0)$. Then fzero iteratively shrinks the interval where fun changes sign to reach a solution.
- 2-element vector — fzero checks that $\text{fun}(x_0(1))$ and $\text{fun}(x_0(2))$ have opposite signs, and errors if they do not. It then iteratively shrinks the interval where fun changes sign to reach a solution. An interval x_0 must be finite; it cannot contain $\pm\text{Inf}$.

Tip Calling fzero with an interval (x_0 with two elements) is often faster than calling it with a scalar x_0 .

Example: 3

Example: [2,17]

Data Types

double

options - Options for solution process

structure, typically created using optimset

Options for solution process, specified as a structure. Create or modify the options structure using optimset. fzero uses these options structure fields.

Display	Level of display: <ul style="list-style-type: none">• 'off' displays no output.• 'iter' displays output at each iteration.• 'final' displays just the final output.• 'notify' (default) displays output only if the function does not converge.
FunValCheck	Check whether objective function values are valid. <ul style="list-style-type: none">• 'on' displays an error when the objective function returns a value that is complex, Inf, or NaN.• The default, 'off', displays no error.
OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none ([]). See “Output Function” on page 9-21.
PlotFcns	Plot various measures of progress while the algorithm executes. Select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]). <ul style="list-style-type: none">• @optimplotx plots the current point.• @optimplotfval plots the function value. For information on writing a custom plot function, see “Plot Functions” on page 9-30.
TolX	Termination tolerance on x, a positive scalar. The default is eps, 2.2204e-16.

Example: `options = optimset('FunValCheck','on')`

Data Types
struct

problem - Root-finding problem

structure

Root-finding problem, specified as a structure with all of the following fields.

<code>objective</code>	Objective function
<code>x0</code>	Initial point for x , scalar or 2-D vector
<code>solver</code>	'fzero'
<code>options</code>	Options structure, typically created using <code>optimset</code>

You can generate `problem` by exporting from Optimization app. See “Importing and Exporting Your Work” on page 5-13 or “Solve Exported Problem” on page 10-151.

Data Types

struct

Output Arguments**`x` - Location of root or sign change**

real scalar

Location of root or sign change, returned as a scalar.

`fval` - Function value at x

real scalar

Function value at x , returned as a scalar.

`exitflag` - Integer encoding the exit condition

integer

Integer encoding the exit condition, meaning the reason `fsolve` stopped its iterations.

- 1 Function converged to a solution x .
- 1 Algorithm was terminated by the output function or plot function.
- 3 NaN or Inf function value was encountered while searching for an interval containing a sign change.
- 4 Complex function value was encountered while searching for an interval containing a sign change.
- 5 Algorithm might have converged to a singular point.
- 6 fzero did not detect a sign change.

output - Information about root-finding process

structure

Information about root-finding process, returned as a structure. The fields of the structure are:

interval	Number of iterations taken to find an interval containing a root
iterations	Number of zero-finding iterations
funcCount	Number of function evaluations
algorithm	'bisection, interpolation'
message	Exit message

Examples

Root Starting From One Point

Calculate π by finding the zero of the sine function near 3.

```
fun = @sin; % function
x0 = 3; % initial point
x = fzero(fun,x0)
```

```
x =
    3.1416
```

Root Starting From an Interval

Find the zero of cosine between 1 and 2.

```
fun = @cos; % function
x0 = [1 2]; % initial interval
x = fzero(fun,x0)
```

```
x =
    1.5708
```

Note that $\cos(1)$ and $\cos(2)$ differ in sign.

Root of a Function Defined by a File

Find a zero of the function $f(x) = x^3 - 2x - 5$.

First, write a file called `f.m`.

```
function y = f(x)
y = x.^3-2*x-5;
```

Save `f.m` on your MATLAB path.

Find the zero of $f(x)$ near 2.

```
fun = @f; % function
x0 = 2; % initial point
z = fzero(fun,x0)
```

```
z =
    2.0946
```

Since $f(x)$ is a polynomial, you can find the same real zero, and a complex conjugate pair of zeros, using the `roots` command.

```
roots([1 0 -2 -5])

ans =
    2.0946
   -1.0473 + 1.1359i
```

```
-1.0473 - 1.1359i
```

Root of Function with Extra Parameter

Find the root of a function that has an extra parameter.

```
myfun = @(x,c) cos(c*x); % parameterized function
c = 2; % parameter
fun = @(x) myfun(x,c); % function of x alone
x = fzero(fun,0.1)
```

```
x =
```

```
0.7854
```

Nondefault Options

Plot the solution process by setting some plot functions.

Define the function and initial point.

```
fun = @(x)sin(cosh(x));
x0 = 1;
```

Examine the solution process by setting options that include plot functions.

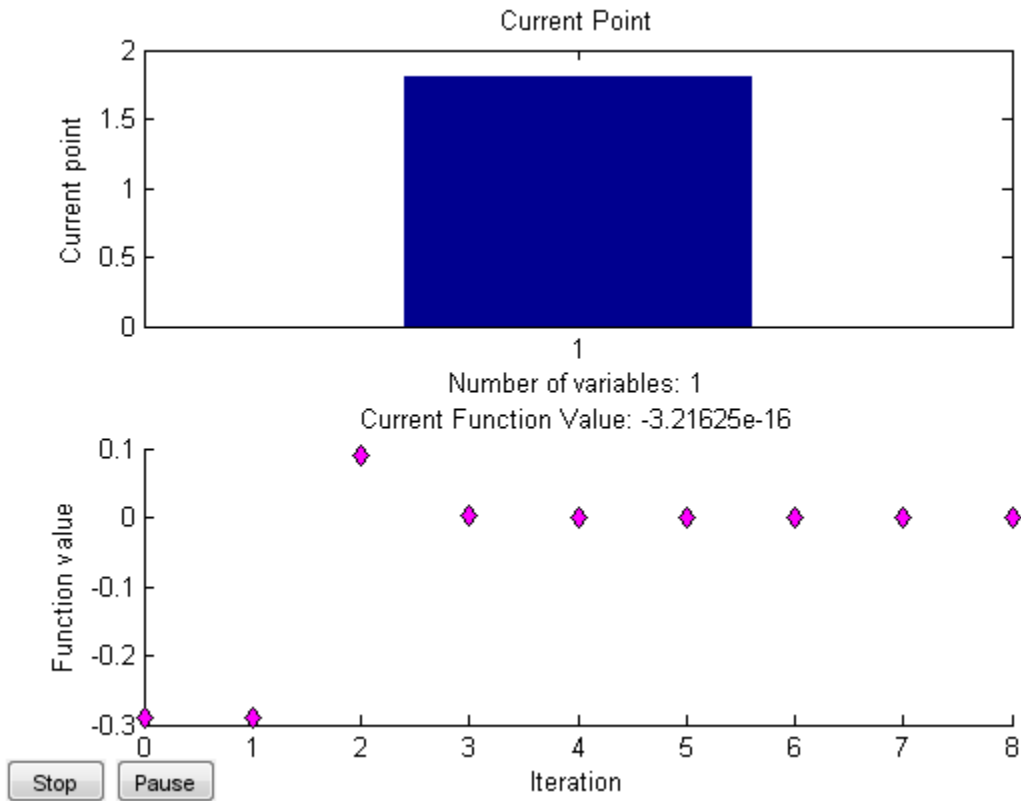
```
options = optimset('PlotFcns',{@optimplotx,@optimplotfval});
```

Run fzero including options.

```
x = fzero(fun,x0,options)
```

```
x =
```

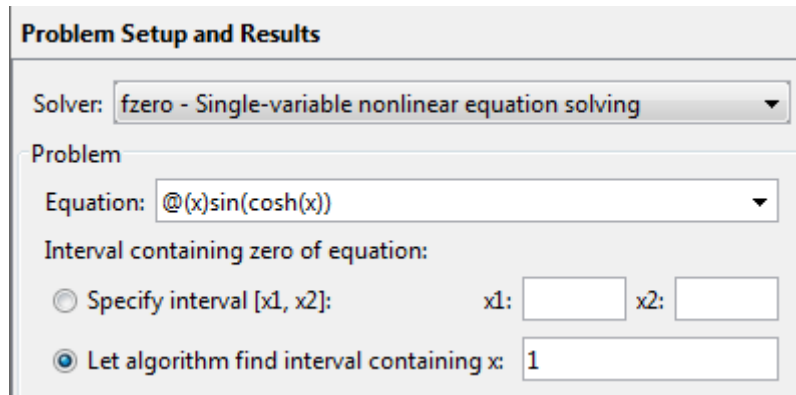
```
1.8115
```



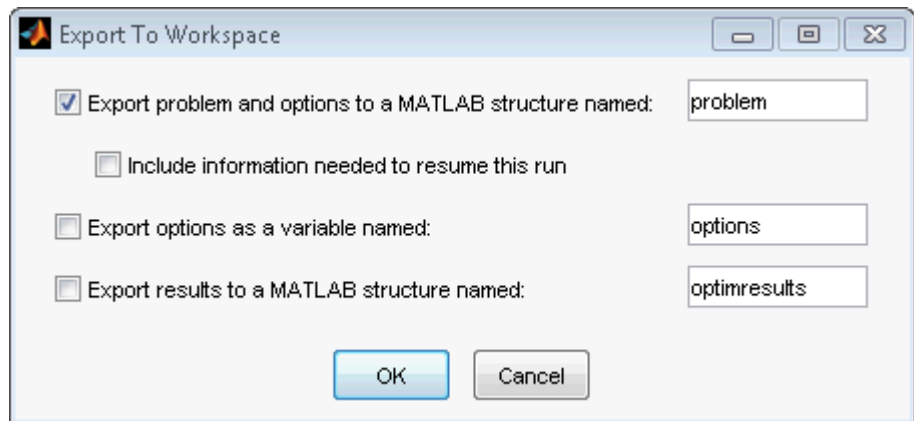
Solve Exported Problem

Solve a problem that is defined by an export from Optimization app.

Define a problem in Optimization app. Enter `optimtool('fzero')`, and fill in the problem as pictured.



Select File > Export to Workspace, and export the problem as pictured to a variable named problem.



Enter the following at the command line.

```
x = fzero(problem)
```

```
x =
```

```
1.8115
```

More Information from Solution

Find the point where $\exp(-\exp(-x)) = x$, and display information about the solution process.

```
fun = @(x) exp(-exp(-x)) - x; % function
x0 = [0 1]; % initial interval
options = optimset('Display','iter'); % show iterations
[x fval exitflag output] = fzero(fun,x0,options)
```

Func-count	x	f(x)	Procedure	
2		1	-0.307799	initial
3	0.544459		0.0153522	interpolation
4	0.566101		0.00070708	interpolation
5	0.567143		-1.40255e-08	interpolation
6	0.567143		1.50013e-12	interpolation
7	0.567143		0	interpolation

```
Zero found in the interval [0, 1]
```

```
x =
```

```
0.5671
```

```
fval =
```

```
0
```

```
exitflag =
```

```
1
```

```
output =
```

```
intervaliterations: 0
```

```
iterations: 5
```

```
funcCount: 7
```

```
algorithm: 'bisection, interpolation'
```

```
message: 'Zero found in the interval [0, 1]'
```

fval = 0 means $\text{fun}(x) = 0$, as desired.

Algorithms

The `fzero` command is a function file. The algorithm, created by T. Dekker, uses a combination of bisection, secant, and inverse quadratic interpolation methods. An Algol 60 version, with some improvements, is given in [1]. A Fortran version, upon which `fzero` is based, is in [2].

References

[1] Brent, R., *Algorithms for Minimization Without Derivatives*, Prentice-Hall, 1973.

[2] Forsythe, G. E., M. A. Malcolm, and C. B. Moler, *Computer Methods for Mathematical Computations*, Prentice-Hall, 1976.

See Also

fminbnd | fsolve | optimset | optimtool | roots

Related Examples

- “Roots of Scalar Functions”

Concepts

- “Passing Extra Parameters” on page 2-53

Purpose

Multiplication with fundamental nullspace basis

Syntax

```
W = fzmult(A,V)
W = fzmult(A,V,'transpose')
[W,L,U,pcol,P] = fzmult(A,V)
W = fzmult(A,V,transpose,L,U,pcol,P)
```

Description

`W = fzmult(A,V)` computes the product W of matrix Z with matrix V , that is, $W = Z*V$, where Z is a fundamental basis for the nullspace of matrix A . A must be a sparse m -by- n matrix where $m < n$, $\text{rank}(A) = m$, and $\text{rank}(A(1:m,1:m)) = m$. V must be p -by- q , where $p = n - m$. If V is sparse W is sparse, else W is full.

`W = fzmult(A,V,'transpose')` computes the product of the transpose of the fundamental basis times V , that is, $W = Z'*V$. V must be p -by- q , where $q = n - m$. `fzmult(A,V)` is the same as `fzmult(A,V,[])`.

`[W,L,U,pcol,P] = fzmult(A,V)` returns the sparse LU-factorization of matrix $A(1:m,1:m)$, that is, $A1 = A(1:m,1:m)$ and $P*A1(:,pcol) = L*U$.

`W = fzmult(A,V,transpose,L,U,pcol,P)` uses the precomputed sparse LU factorization of matrix $A(1:m,1:m)$, that is, $A1 = A(1:m,1:m)$ and $P*A1(:,pcol) = L*U$. `transpose` is either `'transpose'` or `[]`.

The nullspace basis matrix Z is not formed explicitly. An implicit representation is used based on the sparse LU factorization of $A(1:m,1:m)$.

gangstr

Purpose Zero out “small” entries subject to structural rank

Syntax `A = gangstr(M,tol)`

Description `A = gangstr(M,tol)` creates matrix `A` of full structural rank such that `A` is `M` except that elements of `M` that are relatively “small,” based on `tol`, are zeros in `A`. The algorithm decreases `tol`, if needed, until `sprank(A) = sprank(M)`. `M` must have at least as many columns as rows. Default `tol` is `1e-2`.

`gangstr` identifies elements of `M` that are relatively less than `tol` by first normalizing all the rows of `M` to have norm 1. It then examines nonzeros in `M` in a columnwise fashion, replacing with zeros those elements with values of magnitude less than `tol` times the maximum absolute value in that column.

See Also `sprank` | `spy`

Purpose Mixed-integer linear programming (MILP)

Syntax

```
x = intlinprog(f,intcon,A,b)
x = intlinprog(f,intcon,A,b,Aeq,beq)
x = intlinprog(f,intcon,A,b,Aeq,beq,lb,ub)
x = intlinprog(f,intcon,A,b,Aeq,beq,lb,ub,options)
x = intlinprog(problem)
[x,fval,exitflag,output] = intlinprog( ___ )
```

Description Finds the minimum of a problem specified by

$$\min_x f^T x \text{ subject to } \begin{cases} x(\text{intcon}) \text{ are integers} \\ A \cdot x \leq b \\ Aeq \cdot x = beq \\ lb \leq x \leq ub. \end{cases}$$

f , x , intcon , b , beq , lb , and ub are vectors, and A and Aeq are matrices.

You can specify f , intcon , lb , and ub as vectors or arrays. See “Matrix Arguments” on page 2-32.

$x = \text{intlinprog}(f,\text{intcon},A,b)$ solves $\min f^T x$ such that the components of x in intcon are integers, and $A \cdot x \leq b$.

$x = \text{intlinprog}(f,\text{intcon},A,b,Aeq,beq)$ solves the problem above while additionally satisfying the equality constraints $Aeq \cdot x = beq$. Set $A = []$ and $b = []$ if no inequalities exist.

$x = \text{intlinprog}(f,\text{intcon},A,b,Aeq,beq,lb,ub)$ defines a set of lower and upper bounds on the design variables, x , so that the solution is always in the range $lb \leq x \leq ub$. Set $Aeq = []$ and $beq = []$ if no equalities exist.

`x = intlinprog(f,intcon,A,b,Aeq,beq,lb,ub,options)` minimizes using the optimization options specified in `options`. Use `optimoptions` to set these options. Set `lb = []` and `ub = []` if no bounds exist.

`x = intlinprog(problem)` uses a problem structure to encapsulate all solver inputs.

`[x,fval,exitflag,output] = intlinprog(___)`, for any input arguments described above, returns `fval = f'*x`, a value `exitflag` describing the exit condition, and a structure `output` containing information about the optimization process.

Input Arguments

f - Coefficient vector

real vector

Coefficient vector, specified as a vector of doubles representing the objective function, $f' * x$. The notation assumes that `f` is a column vector, but you are free to use a row vector.

`f` can also be an array. Internally, `intlinprog` converts an array `f` to the vector `f(:)`.

If you specify `f = []`, `intlinprog` tries to find a feasible point without trying to minimize an objective function.

Example: `f = [4;2;-1.7];`

Data Types

double

intcon - Vector of integer constraints

vector of integers

Vector of integer constraints, specified as a vector of positive integers. The values in `intcon` indicate the components of the decision variable `x` that are integer-valued. `intcon` has values from 1 through `numel(f)`.

`intcon` can also be an array. Internally, `intlinprog` converts an array `intcon` to the vector `intcon(:)`.

Example: `intcon = [1,2,7]` means $x(1)$, $x(2)$, and $x(7)$ take only integer values.

Data Types

double

A - Linear inequality constraint matrix

real matrix

Linear inequality constraint matrix, specified as a matrix of doubles. A represents the linear coefficients in the constraints $A*x \leq b$. A has size M -by- N , where M is the number of constraints and $N = \text{numel}(f)$. To save memory, A can be sparse.

Example: `A = [4,3;2,0;4,-1]`; means three linear inequalities (three rows) for two decision variables (two columns).

Data Types

double

b - Linear inequality constraint vector

real vector

Linear inequality constraint vector, specified as a vector of doubles. b represents the constant vector in the constraints $A*x \leq b$. b has length M , where A is M -by- N .

Example: `[4,0]`

Data Types

double

Aeq - Linear equality constraint matrix

`[]` (default) | real matrix

Linear equality constraint matrix, specified as a matrix of doubles. A_{eq} represents the linear coefficients in the constraints $A_{eq}*x = b_{eq}$. A_{eq} has size M_{eq} -by- N , where M_{eq} is the number of constraints and $N = \text{numel}(f)$. To save memory, A_{eq} can be sparse.

Example: $A = [4, 3; 2, 0; 4, -1]$; means three linear inequalities (three rows) for two decision variables (two columns).

Data Types

double

beq - Linear equality constraint vector

[] (default) | real vector

Linear equality constraint vector, specified as a vector of doubles. `beq` represents the constant vector in the constraints $A_{eq}x = beq$. `beq` has length `Meq`, where `Aeq` is `Meq`-by-`N`.

Example: $[4, 0]$

Data Types

double

lb - Lower bounds

[] (default) | real vector or array

Lower bounds, specified as a vector or array of doubles. `lb` represents the lower bounds elementwise in $lb \leq x \leq ub$.

Internally, `intlinprog` converts an array `lb` to the vector `lb(:)`.

Example: $lb = [0; -Inf; 4]$ means $x(1) \geq 0$, $x(3) \geq 4$.

Data Types

double

ub - Upper bounds

[] (default) | real vector or array

Upper bounds, specified as a vector or array of doubles. `ub` represents the upper bounds elementwise in $lb \leq x \leq ub$.

Internally, `intlinprog` converts an array `ub` to the vector `ub(:)`.

Example: $ub = [Inf; 4; 10]$ means $x(2) \leq 4$, $x(3) \leq 10$.

Data Types

double

options - Options for intlinprog

options created using `optimoptions`

Options for `intlinprog`, specified as the output of `optimoptions`.

Option	Description	Default
BranchingRule	<p>Rule for choosing the component for branching:</p> <ul style="list-style-type: none"> 'maxpscost' — The fractional component with maximum pseudocost. See “Branch and Bound” on page 6-161. 'mostfractional' — The component whose fractional part is closest to 1/2. 'maxfun' — The fractional component with maximal corresponding component in the absolute value of objective vector f. 	'maxpscost'
CutGeneration	<p>Level of cut generation (see “Cut Generation” on page 6-160):</p> <ul style="list-style-type: none"> 'none' — No cuts. Makes <code>CutGenerationMaxIter</code> irrelevant. 'basic' — Normal cut generation. 'intermediate' — Use more cut types. 'advanced' — Use most cut types. 	'basic'
CutGenMaxIter	<p>Number of passes through all cut generation methods before entering the branch-and-bound phase, an integer from 1 through 50. Disable cut generation by setting the <code>CutGeneration</code> option to 'none'.</p>	10

Option	Description	Default
Display	Level of display (see “Iterative Display” on page 3-17): <ul style="list-style-type: none">• 'off' or 'none' — No iterative display• 'final' — Show final values only• 'iter' — Show iterative display	'iter'
Heuristics	Algorithm for searching for feasible points (see “Heuristics for Finding Feasible Solutions” on page 6-161): <ul style="list-style-type: none">• 'none'• 'rss'• 'round'• 'rins'	'rss'
HeuristicsMaxNodes	MaxNodes, positive integer that bounds the number of nodes intlinprog can explore in its branch-and-bound search for feasible points. See “Heuristics for Finding Feasible Solutions” on page 6-161.	50

Option	Description	Default
IPPreprocess	<p>Types of integer preprocessing (see “Mixed-Integer Program Preprocessing” on page 6-159):</p> <ul style="list-style-type: none"> • 'none' — Use very few integer preprocessing steps. • 'basic' — Use a moderate number of integer preprocessing steps. • 'advanced' — Use all available integer preprocessing steps. 	'basic'
LPMaxIter	Strictly positive integer, the maximum number of simplex algorithm iterations per node during the branch-and-bound process.	3e4
LPPreprocess	<p>Type of preprocessing for the solution to the relaxed linear program (see “Linear Program Preprocessing” on page 6-158):</p> <ul style="list-style-type: none"> • 'none' — No preprocessing. • 'basic' — Use preprocessing. 	'basic'
MaxNodes	Strictly positive integer that is the maximum number of nodes intlinprog explores in its branch-and-bound process.	1e7
MaxNumFeasPoints	Strictly positive integer. intlinprog stops if it finds MaxNumFeasPoints integer feasible points.	Inf
MaxTime	Positive real that is the maximum time in seconds that intlinprog runs.	7200

Option	Description	Default
NodeSelection	<p>Choose the node to explore next.</p> <ul style="list-style-type: none"> 'simplebestproj' — Best projection. See “Branch and Bound” on page 6-161. 'minobj' — Explore the node with the minimum objective function. 'mininfeas' — Explore the node with the minimal sum of integer infeasibilities. See “Branch and Bound” on page 6-161. 	'simplebestproj'
ObjectiveCutoff	<p>Real greater than -Inf. During the branch-and-bound calculation, intlinprog discards any node where the linear programming solution has an objective value exceeding ObjectiveCutoff.</p>	Inf
RelObjThreshold	<p>Nonnegative real. intlinprog changes the current feasible solution only when it locates another with an objective function value that is at least RelObjThreshold lower: $(\text{fold} - \text{fnew}) / (1 + \text{fold}) > \text{RelObjThreshold}$.</p>	1e-4
RootLPAlgorithm	<p>Algorithm for solving linear programs:</p> <ul style="list-style-type: none"> 'dual-simplex' — Dual simplex algorithm 'primal-simplex' — Primal simplex algorithm 	'dual-simplex'
RootLPMaxIter	<p>Nonnegative integer that is the maximum number of simplex algorithm iterations to solve the initial linear programming problem.</p>	3e4

Option	Description	Default
TolCon	Real from 1e-9 through 1e-3 that is the maximum discrepancy that linear constraints can have and still be considered satisfied. TolCon is not a stopping criterion.	1e-4
TolFunLP	Nonnegative real where reduced costs must exceed TolFunLP for a variable to be taken into the basis.	1e-7
TolGapAbs	Nonnegative real. intlinprog stops if the difference between the internally calculated upper (U) and lower (L) bounds on the objective function is less than or equal to TolGapAbs: $U - L \leq \text{TolGapAbs}$.	0
TolGapRel	Real from 0 through 1. intlinprog stops if the relative difference between the internally calculated upper (U) and lower (L) bounds on the objective function is less than or equal to TolGapRel: $(U - L) / (\text{abs}(U) + 1) \leq \text{TolGapRel}$.	1e-4
TolInteger	Real from 1e-6 through 1e-3, where the maximum deviation from integer that a component of the solution x can have and still be considered an integer. TolInteger is not a stopping criterion.	1e-5

Example: `options = optimoptions('intlinprog','MaxTime',120)`

problem - Structure encapsulating inputs and options

structure

Structure encapsulating the inputs and options, specified with the following fields.

intlinprog

<code>f</code>	Vector representing objective $f'x$ (required)
<code>intcon</code>	Vector indicating variables that take integer values (required)
<code>Aineq</code>	Matrix in linear inequality constraints $A_{ineq}x \leq b_{ineq}$
<code>bineq</code>	Vector in linear inequality constraints $A_{ineq}x \leq b_{ineq}$
<code>Aeq</code>	Matrix in linear equality constraints $A_{eq}x = b_{eq}$
<code>b eq</code>	Vector in linear equality constraints $A_{eq}x = b_{eq}$
<code>lb</code>	Vector of lower bounds
<code>ub</code>	Vector of upper bounds
<code>solver</code>	'intlinprog' (required)
<code>options</code>	Options created using <code>optimoptions</code> (required)

You must specify at least these fields in the problem structure. Other fields are optional:

- `f`
- `intcon`
- `solver`
- `options`

```
Example: problem.f = [1,2,3];  
problem.intcon = [2,3];  
problem.options = optimoptions('intlinprog');  
problem.Aineq = [-3,-2,-1];  
problem.bineq = -20;  
problem.lb = [-6.1,-1.2,7.3];  
problem.solver = 'intlinprog';
```

Data Types
struct

Output Arguments**x - Solution**

real vector

Solution, returned as a vector that minimizes $f' * x$ subject to all bounds, integer constraints, and linear constraints.

When a problem is infeasible or unbounded, x is $[]$.

fval - Objective value

real scalar

Objective value, returned as the scalar value $f' * x$ at the solution x .

When a problem is infeasible or unbounded, $fval$ is $[]$.

exitflag - Algorithm stopping condition

integer

Algorithm stopping condition, returned as an integer identifying the reason the algorithm stopped. The following lists the values of `exitflag` and the corresponding reasons `intlinprog` stopped.

2	<code>intlinprog</code> stopped prematurely. Integer feasible point found.
1	<code>intlinprog</code> converged to the solution x .
0	<code>intlinprog</code> stopped prematurely. No integer feasible point found.
-2	No feasible point found.
-3	Root LP problem is unbounded.

The exit message can give more detailed information on the reason `intlinprog` stopped, such as exceeding a tolerance.

output - Solution process summary

structure

Solution process summary, returned as a structure containing information about the optimization process.

intlinprog

relativegap	Relative difference between upper (U) and lower (L) bounds of the objective function that intlinprog calculates in its branch-and-bound algorithm. $\text{relativegap} = (U - L) / (\text{abs}(U) + 1)$ If intcon = [], relativegap = [].
absolutegap	Difference between upper and lower bounds of the objective function that intlinprog calculates in its branch-and-bound algorithm. If intcon = [], absolutegap = [].
numfeaspoints	Number of integer feasible points found. If intcon = [], numfeaspoints = []. Also, if the initial relaxed problem is infeasible, numfeaspoints = [].
numnodes	Number of nodes in branch-and-bound algorithm. If the solution was found during preprocessing or during the initial cuts, numnodes = 0. If intcon = [], numnodes = [].
constrviolation	Constraint violation that is positive for violated constraints. $\text{constrviolation} = \max([0; \text{norm}(\text{Aeq} \cdot x - \text{beq}, \text{inf}); (\text{lb} - x); (x - \text{ub}); (\text{Ai} \cdot x - \text{bi})])$
message	Exit message.

Examples

Solve an MILP with Linear Inequalities

Solve the problem

$$\min_x 8x_1 + x_2 \text{ subject to } \begin{cases} x_2 \text{ is an integer} \\ x_1 + 2x_2 \geq -14 \\ -4x_1 - x_2 \leq -33 \\ 2x_1 + x_2 \leq 20. \end{cases}$$

Write the objective function vector and vector of integer variables.

```
f = [8;1];
intcon = 2;
```

Convert all inequalities into the form $A*x \leq b$ by multiplying “greater than” inequalities by -1.

```
A = [-1,-2;
      -4,-1;
      2,1];
b = [14;-33;20];
```

Call intlinprog.

```
x = intlinprog(f,intcon,A,b)
```

```
LP: Optimal objective value is 59.000000.
```

Optimal solution found.

Intlinprog stopped at the root node because the objective value is within the tolerance. The tolerance is specified by options.TolGapAbs = 0 (the default value). The intcon variables are integer. The tolerance for integer variables is specified by options.TolInteger = 1e-05 (the default value).

```
x =
    6.5000
    7.0000
```

Solve an MILP with All Types of Constraints

Solve the problem

$$\min_x (-3x_1 - 2x_2 - x_3) \text{ subject to } \begin{cases} x_3 \text{ binary} \\ x_1, x_2 \geq 0 \\ x_1 + x_2 + x_3 \leq 7 \\ 4x_1 + 2x_2 + x_3 = 12 \end{cases}$$

Write the objective function vector and vector of integer variables.

```
f = [-3; -2; -1];  
intcon = 3;
```

Write the linear inequality constraints.

```
A = [1, 1, 1];  
b = 7;
```

Write the linear equality constraints.

```
Aeq = [4, 2, 1];  
beq = 12;
```

Write the bound constraints.

```
lb = zeros(3, 1);  
ub = [Inf; Inf; 1]; % Enforces x(3) is binary
```

Call `intlinprog`.

```
x = intlinprog(f, intcon, A, b, Aeq, beq, lb, ub)
```

```
LP: Optimal objective value is -12.000000.
```

```
Optimal solution found.
```


Intlinprog stopped at the root node because the objective value is within the tolerance. The options.TolGapAbs = 0 (the default value). The intcon variables are integer. The options.TolInteger = 1e-05 (the default value).

```
x =
     0
    5.5000
    1.0000
```

Solve an MILP with Nondefault Options

Solve the problem

$$\min_x (-3x_1 - 2x_2 - x_3) \text{ subject to } \begin{cases} x_3 \text{ binary} \\ x_1, x_2 \geq 0 \\ x_1 + x_2 + x_3 \leq 7 \\ 4x_1 + 2x_2 + x_3 = 12 \end{cases}$$

without showing iterative display.

Specify the solver inputs.

```
f = [-3;-2;-1];
intcon = 3;
A = [1,1,1];
b = 7;
Aeq = [4,2,1];
beq = 12;
lb = zeros(3,1);
ub = [Inf;Inf;1]; % enforces x(3) is binary
```

Specify no display.

```
options = optimoptions('intlinprog','Display','off');
```

Run the solver.

```
x = intlinprog(f,intcon,A,b,Aeq,beq,lb,ub,options)
```

```
x =
```

```
    0  
  5.5000  
  1.0000
```

Use a Problem Structure

Solve the problem

$$\min_x (-3x_1 - 2x_2 - x_3) \text{ subject to } \begin{cases} x_3 \text{ binary} \\ x_1, x_2 \geq 0 \\ x_1 + x_2 + x_3 \leq 7 \\ 4x_1 + 2x_2 + x_3 = 12 \end{cases}$$

using iterative display. Use a problem structure as the intlinprog input.

Specify the solver inputs.

```
f = [-3;-2;-1];  
intcon = 3;  
A = [1,1,1];  
b = 7;  
Aeq = [4,2,1];  
beq = 12;  
lb = zeros(3,1);  
ub = [Inf;Inf;1]; % enforces x(3) is binary  
options = optimoptions('intlinprog','Display','off');
```

Insert the inputs into a problem structure. Include the solver name.

```
problem = struct('f',f,'intcon',intcon,...  
    'Aineq',A,'bineq',b,'Aeq',Aeq,'beq',beq,...  
    'lb',lb,'ub',ub,'options',options,...  
    'solver','intlinprog');
```

Run the solver.

```
x = intlinprog(problem)
```

```
x =
```

```

         0
    5.5000
    1.0000
```

Examine the MILP Solution and Process

Call `intlinprog` with more outputs to see solution details and process.

The goal is to solve the problem

$$\min_x (-3x_1 - 2x_2 - x_3) \text{ subject to } \begin{cases} x_3 \text{ binary} \\ x_1, x_2 \geq 0 \\ x_1 + x_2 + x_3 \leq 7 \\ 4x_1 + 2x_2 + x_3 = 12 \end{cases}$$

Specify the solver inputs.

```
f = [-3;-2;-1];
intcon = 3;
A = [1,1,1];
b = 7;
Aeq = [4,2,1];
beq = 12;
lb = zeros(3,1);
ub = [Inf;Inf;1]; % enforces x(3) is binary
```

Call `intlinprog` with all outputs.

```
[x,fval,exitflag,output] = intlinprog(f,intcon,A,b,Aeq,beq,lb,ub)
```

```
LP:                Optimal objective value is -12.000000.
```

intlinprog

Optimal solution found.

Intlinprog stopped at the root node because the objective value is within options.TolGapAbs = 0 (the default value). The intcon variables are integer options.TolInteger = 1e-05 (the default value).

x =

```
      0
    5.5000
    1.0000
```

fval =

```
-12.0000
```

exitflag =

```
1
```

output =

```
    relativegap: 0
    absolutegap: 0
    numfeaspoints: 1
    numnodes: 0
    constrviolation: 1.7764e-15
    message: 'Optimal solution found.'
```

Intlinprog stopped at the root node because the objective v...'

The output structure shows `numnodes` is 0. This means `intlinprog` solved the problem before branching. This is one indication that the result is reliable. Also, the `absolutegap` and `relativegap` fields are 0. This is another indication that the result is reliable.

Tips

- To specify binary variables, set the variables to be integers in `intcon`, and give them lower bounds of 0 and upper bounds of 1.
- Save memory by specifying sparse linear constraint matrices `A` and `Aeq`. However, you cannot use sparse matrices for `b` and `beq`.
- To provide logical indices for integer components, meaning a binary vector with 1 indicating an integer, convert to `intcon` form using `find`. For example,

```
logicalindices = [1,0,0,1,1,0,0];
intcon = find(logicalindices)
```

```
intcon =
```

```
     1     4     5
```

Limitations

- Often, some supposedly integer-valued components of the solution `x(intcon)` are not precisely integers. `intlinprog` deems as integers all solution values within the `TolInteger` tolerance of an integer.

To round all supposed integers to be exactly integers, use the `round` function.

```
x(intcon) = round(x(intcon));
```

Caution

Rounding solutions can cause the solution to become infeasible. Check feasibility after rounding:

```
max(A*x - b) % See if entries are not too positive, so have small infeasibility
max(abs(Aeq*x - beq)) % See if entries are near enough to zero
max(x - ub) % Positive entries are violated bounds
max(lb - x) % Positive entries are violated bounds
```

- `intlinprog` does not enforce that solution components be integer-valued when their absolute values exceed $2.1e9$. When your solution has such components, `intlinprog` warns you. If you receive this warning, check the solution to see whether supposedly integer-valued components of the solution are close to integers.
- `intlinprog` does not allow components of the problem, such as coefficients in `f`, `A`, or `ub`, to exceed $1e25$ in absolute value. If you try to run `intlinprog` with such a problem, `intlinprog` issues an error.
- Currently, you cannot run `intlinprog` in the Optimization app.

Definitions

MILP

Mixed-integer linear programming definition.

MILP means find the minimum of a problem specified by

$$\min_x f^T x \text{ subject to } \begin{cases} x(\text{intcon}) \text{ are integers} \\ A \cdot x \leq b \\ Aeq \cdot x = beq \\ lb \leq x \leq ub. \end{cases}$$

f , x , intcon , b , beq , lb , and ub are vectors, and A and Aeq are matrices.

You can specify f , intcon , lb , and ub as vectors or arrays. See “Matrix Arguments” on page 2-32.

See Also

`linprog` | `optimoptions`

Related Examples

- “Mixed-Integer Linear Programming Basics” on page 6-169
- “Factory, Warehouse, Sales Allocation Model” on page 6-173
- “Travelling Salesman Problem” on page 6-185
- “Solve Sudoku Puzzles Via Integer Programming” on page 6-192

Concepts

- “Mixed-Integer Linear Programming Algorithms” on page 6-157
- “Tuning Integer Linear Programming” on page 6-166
- “Optimization Problem Setup”

Purpose

Find minimum of constrained or unconstrained nonlinear multivariable function using KNITRO third-party libraries

Equation

Finds the minimum of a problem specified by

$$\min_x f(x) \text{ such that } \begin{cases} c(x) \leq 0 \\ ceq(x) = 0 \\ A \cdot x \leq b \\ Aeq \cdot x = beq \\ lb \leq x \leq ub, \end{cases}$$

where b and beq are vectors, A and Aeq are matrices, $c(x)$ and $ceq(x)$ are functions that return vectors, and $f(x)$ is a function that returns a scalar. $f(x)$, $c(x)$, and $ceq(x)$ can be nonlinear functions. All constraints are optional; `ktrlink` can minimize unconstrained problems.

x , lb , and ub can be passed as vectors or matrices; see “Matrix Arguments” on page 2-32.

Note To use `ktrlink`, you must have KNITRO third-party libraries, and have a path that includes these libraries. See “`ktrlink`: An Interface to KNITRO Libraries” on page 8-2.

Syntax

```
x = ktrlink(fun,x0)
x = ktrlink(fun,x0,A,b)
x = ktrlink(fun,x0,A,b,Aeq,beq)
x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub)
x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon)
x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)
x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options,
knitroOptions)
[x,fval] = ktrlink(...)
[x,fval,exitflag] = ktrlink(...)
[x,fval,exitflag,output] = ktrlink(...)
```



```
[x,fval,exitflag,output,lambda] = ktrlink(...)
```

Description

ktrlink attempts to find a minimum of a scalar function of several variables starting at an initial estimate. This is generally referred to as *constrained or unconstrained nonlinear optimization*, or *nonlinear programming*.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the objective function and nonlinear constraint functions, if necessary.

`x = ktrlink(fun,x0)` starts at `x0` and attempts to find a minimizer `x` of the function described in `fun`, subject to no constraints. `x0` can be a scalar, vector, or matrix.

`x = ktrlink(fun,x0,A,b)` minimizes `fun` subject to the linear inequalities $A*x \leq b$.

`x = ktrlink(fun,x0,A,b,Aeq,beq)` minimizes `fun` subject to the linear equalities $Aeq*x = beq$ as well as $A*x \leq b$. If no inequalities exist, set `A = []` and `b = []`.

`x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub)` defines a set of lower and upper bounds on the design variables in `x`, so that the solution is always in the range $lb \leq x \leq ub$. If no equalities exist, set `Aeq = []` and `beq = []`. If `x(i)` is unbounded below, set `lb(i) = -Inf`, and if `x(i)` is unbounded above, set `ub(i) = Inf`.

`x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon)` subjects the minimization to the nonlinear inequalities `c(x)` and the equalities `ceq(x)` defined in `nonlcon`. `fmincon` optimizes such that $c(x) \leq 0$ and $ceq(x) = 0$. If no bounds exist, set `lb = []` and/or `ub = []`.

`x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options. If there are no nonlinear inequality or equality constraints, set `nonlcon = []`.

`x = ktrlink(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options, knitroOptions)` minimizes with the KNITRO options specified in the text file `knitroOptions`. All options given in `options` are ignored except for `HessFcn`, `HessMult`, `HessPattern`, and `JacobPattern`.

`[x,fval] = ktrlink(...)` returns the value of the objective function `fun` at the solution `x`.

`[x,fval,exitflag] = ktrlink(...)` returns `exitflag`, which describes the exit condition of the KNITRO solver.

`[x,fval,exitflag,output] = ktrlink(...)` returns a structure `output` with information about the optimization.

`[x,fval,exitflag,output,lambda] = ktrlink(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the output `fval` is `[]`.

Components of `x0` that violate the bounds `lb` `x` `ub` are reset to the interior of the box defined by the bounds. Components that respect the bounds are not changed.

Input Arguments

“Function Arguments” on page 9-2 contains descriptions of arguments passed to `ktrlink`. “Options” on page 10-52 provides the function-specific details for the options values. This section provides function-specific details for `fun` and `nonlcon`.

fun The function to be minimized. `fun` is a function that accepts a vector `x` and returns a scalar `f`, the objective function evaluated at `x`. `fun` can be specified as a function handle for a function file:

```
x = ktrlink(@myfun,x0,A,b)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...                    % Compute function value at x
```

`fun` can also be a function handle for an anonymous function.

```
x = ktrlink(@(x)norm(x)^2,x0,A,b);
```

If you can compute the gradient of `fun` *and* the `GradObj` option is 'on', as set by

```
options = optimoptions('ktrlink','GradObj','on')
```

then `fun` must return the gradient vector `g(x)` in the second output argument.

If you can compute the Hessian matrix, there are several ways to pass the Hessian to `ktrlink`. See “Hessian” on page 10-183 for details.

nonlcon The function that computes the nonlinear inequality constraints $c(x) \leq 0$ and the nonlinear equality constraints $ceq(x) = 0$. `nonlcon` accepts a vector `x` and returns the two vectors `c` and `ceq`. `c` contains the nonlinear inequalities evaluated at `x`, and `ceq` contains the nonlinear equalities evaluated at `x`. The function `nonlcon` can be specified as a function handle.

```
x = ktrlink(@myfun,x0,A,b,Aeq,beq,lb,ub,@mycon)
```

where `mycon` is a MATLAB function such as

```
function [c,ceq] = mycon(x)
c = ...                    % Compute nonlinear inequalities at x.
```

```
ceq = ... % Compute nonlinear equalities at x.
```

If you can compute the gradients of the constraints *and* the GradConstr option is 'on', as set by

```
options = optimoptions('ktrlink','GradConstr','on')
```

then nonlcon must also return GC, the gradient of $c(x)$, and GCEq, the gradient of $ceq(x)$, in the third and fourth output arguments respectively. See “Nonlinear Constraints” on page 2-37 for details.

Note Because Optimization Toolbox functions only accept inputs of type `double`, user-supplied objective and nonlinear constraint functions must return outputs of type `double`.

“Passing Extra Parameters” on page 2-53 explains how to parameterize the nonlinear constraint function `nonlcon`, if necessary.

Output Arguments

“Function Arguments” on page 9-2 contains descriptions of arguments returned by `ktrlink`. This section provides function-specific details for `exitflag`, `lambda`, and output:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. For more information, see the KNITRO documentation at http://www.ziena.com/
<code>lambda</code>	Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields of the structure are
	<code>lower</code> Lower bounds <code>lb</code>
	<code>upper</code> Upper bounds <code>ub</code>
	<code>ineqlin</code> Linear inequalities

	<code>eqlin</code>	Linear equalities
	<code>ineqnonlin</code>	Nonlinear inequalities
	<code>eqnonlin</code>	Nonlinear equalities
output		Structure containing information about the optimization. The fields of the structure are:
	<code>iterations</code>	Number of iterations taken
	<code>funcCount</code>	Number of function evaluations
	<code>constrviolation</code>	Maximum of constraint violations (interior-point algorithm only)
	<code>firstorderopt</code>	Measure of first-order optimality

Hessian

`ktrlink` can optionally use a user-supplied Hessian, the matrix of second derivatives of the Lagrangian, namely,

$$\nabla_{xx}^2 L(x, \lambda) = \nabla^2 f(x) + \sum \lambda_i \nabla^2 c_i(x) + \sum \lambda_i \nabla^2 ceq_i(x). \quad (10-2)$$

If you don't supply a Hessian, KNITRO software estimates it.

To provide a Hessian, the syntax is

```
hessian = hessianfcn(x, lambda)
```

`hessian` is an n -by- n matrix, sparse or dense, where n is the number of variables. `lambda` is a structure with the Lagrange multiplier vectors associated with the nonlinear constraints:

```
lambda.ineqnonlin
lambda.eqnonlin
```

KNITRO software computes `lambda.hessianfcn` must calculate the sums in Equation 10-2. Indicate that you are supplying a Hessian by

```
options = optimoptions('ktrlink','Hessian','user-supplied',...  
                      'HessFcn',@hessianfcn);
```

There are several more options for Hessians:

- `options = optimoptions('ktrlink','Hessian','bfgs');`

The KNITRO solver calculates the Hessian by a dense quasi-Newton approximation.

- `options = optimoptions('ktrlink','Hessian',...
{'lbfgs',positive integer});`

The KNITRO solver calculates the Hessian by a limited-memory, large-scale quasi-Newton approximation. The positive integer specifies how many past iterations should be remembered.

- `options =
optimoptions('ktrlink','Hessian','fin-diff-grads',...
'SubproblemAlgorithm','cg','GradObj','on',...
'GradConstr','on');`

The KNITRO solver calculates a Hessian-times-vector product by finite differences of the gradient(s). You must supply the gradient of the objective function, and also gradients of any nonlinear constraint functions.

- `options =
optimoptions('ktrlink','Hessian','user-supplied',...
'SubproblemAlgorithm','cg','HessMult',@HessMultFcn);`

The KNITRO solver uses a Hessian-times-vector product. You must supply the function `HessMultFcn`, which returns an n -by-1 vector. The `HessMult` option enables you to pass the result of multiplying the Hessian by a vector without calculating the Hessian.

The syntax for the 'HessMult' option is:

```
w = HessMultFcn(x,lambda,v);
```

The result w should be the product $H*v$, where H is the Hessian at x , λ is the Lagrange multiplier (computed by KNITRO software), and v is a vector.

Options

Optimization options used by ktrlink. Use `optimoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information. For example:

```
options = optimoptions('ktrlink','Algorithm','active-set');
```

Option	Description
Algorithm	Choose a KNITRO optimization algorithm: 'interior-point' or 'active-set'. Default: 'interior-point'.
AlwaysHonorConstraints	The default 'bounds' ensures that bound constraints are satisfied at every iteration. Disable by setting to 'none'.
Display	Level of display: <ul style="list-style-type: none"> • 'off' or 'none' displays no output. • 'iter' displays output at each iteration, and gives the default exit message. • 'iter-detailed' displays output at each iteration, and gives the technical exit message. • 'notify' displays output only if the function does not converge, and gives the default exit message. • 'notify-detailed' displays output only if the function does not converge, and gives the technical exit message.

Option	Description
	<ul style="list-style-type: none">• 'final' (default) displays just the final output, and gives the default exit message.• 'final-detailed' displays just the final output, and gives the technical exit message.
FinDiffType	Finite differences, used to estimate gradients, are either 'forward' (the default), or 'central' (centered). 'central' takes twice as many function evaluations but should be more accurate. 'central' differences might violate bounds during their evaluation.
FunValCheck	Check whether objective function values are valid. <ul style="list-style-type: none">• 'on' displays an error when the objective function returns a value that is complex, Inf, or NaN.• The default, 'off', displays no error.
GradConstr	Gradient for nonlinear constraint functions defined by the user. When set to 'on', ktrlink expects the constraint function to have four outputs, as described in <code>nonlcon</code> in the “Input Arguments” on page 10-180 section. When set to the default, 'off', gradients of the nonlinear constraints are estimated by finite differences. If you want to use sparse constraint gradients, you must set the sparsity pattern with <code>JacobPattern</code> , as described

Option	Description
	in “Sparsity Pattern for Nonlinear Constraints” on page 8-10.
	<hr/> Caution JacobPattern has transposed orientation compared to the constraint gradients. <hr/>
GradObj	Gradient for the objective function defined by the user. See the preceding description of <code>fun</code> to see how to define the gradient in <code>fun</code> . Set to 'on' to have KNITRO use a user-defined gradient of the objective function. The default 'off' causes KNITRO to estimate gradients using finite differences.
HessFcn	Function handle to a user-supplied Hessian (see “Hessian” on page 10-183). Default: [].
Hessian	Chooses how ktrlink calculates the Hessian (see “Hessian” on page 10-183). Default: 'bfgs'.
HessMult	Handle to a user-supplied function that gives a Hessian-times-vector product (see “Hessian” on page 10-183). Default: [].
HessPattern	Sparsity pattern of objective function Hessian (see “Hessian” on page 10-183). Default: 'sparse(ones(numberOfVariables))'.
InitBarrierParam	Initial barrier value. A value above the default 0.1 might help, especially if the objective or constraint functions are large.

Option	Description
<code>InitTrustRegionRadius</code>	Initial radius of the trust region. On badly-scaled problems choose a value smaller than the default, \sqrt{n} , where n is the number of variables.
<code>JacobPattern</code>	<p>Sparsity pattern of the Jacobian of the nonlinear constraint matrix for finite differencing or for a gradient function in <code>nonlcon</code>.</p> <p>Even when you include the nonlinear constraint gradient in <code>nonlcon</code>, if you want <code>ktrlink</code> to use sparse gradients, you must pass the sparsity pattern via <code>JacobPattern</code>, as described in “Sparsity Pattern for Nonlinear Constraints” on page 8-10. Default: <code>'sparse(ones(Jrows,Jcols))'</code>.</p>
<code>MaxIter</code>	Maximum number of iterations allowed. Default: 10000.
<code>MaxProjCGIter</code>	A tolerance (stopping criterion) for the number of projected conjugate gradient iterations; this is an inner iteration, not the number of iterations of the algorithm. Default: <code>'2*(numberOfVariables-numberOfEqualities)'</code> .
<code>ObjectiveLimit</code>	A tolerance (stopping criterion). If the objective function value goes below <code>ObjectiveLimit</code> and the iterate is feasible, the iterations halt, since the problem is presumably unbounded. Default: <code>-1e20</code> .

Option	Description
ScaleProblem	The default 'obj-and-constr' causes the algorithm to normalize all constraints and the objective function. Disable by setting to 'none'.
SubproblemAlgorithm	Determines how the iteration step is calculated. The default 'ldl-factorization' is usually faster than 'cg' (conjugate gradient), though 'cg' may be faster for large problems with dense Hessians.
TolCon	Termination tolerance on the constraint violation. Default: 1e-6.
TolFun	Termination tolerance on the function value. Default: 1e-6.
TolX	Termination tolerance on x. Default: 1e-15.

KNITRO® Options

You can set options for the KNITRO libraries and pass them in a text file. The text file should consist of lines of text with the name of an option followed by blank space and then the desired value of the option. For example, to select the maximum run time to be less than 100 seconds, and to use an adaptive algorithm for changing the multiplier μ , create a text file containing the following lines:

```
ms_maxtime_real 100
bar_murule adaptive
```

For full details about the structure of the file and all possible options, see the KNITRO documentation at <http://www.ziena.com/>.

References

[1] <http://www.ziena.com/>

See Also

fminbnd | fmincon | fminsearch | fminunc | optimoptions

How To

- `function_handle`
- “ktrlink: An Interface to KNITRO Libraries” on page 8-2

Purpose Solve linear programming problems

Equation Finds the minimum of a problem specified by

$$\min_x f^T x \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ lb \leq x \leq ub. \end{cases}$$

f , x , b , beq , lb , and ub are vectors, and A and Aeq are matrices.

Syntax

```
x = linprog(f,A,b)
x = linprog(f,A,b,Aeq,beq)
x = linprog(f,A,b,Aeq,beq,lb,ub)
x = linprog(f,A,b,Aeq,beq,lb,ub,x0)
x = linprog(f,A,b,Aeq,beq,lb,ub,x0,options)
x = linprog(problem)
[x,fval] = linprog(...)
[x,fval,exitflag] = linprog(...)
[x,fval,exitflag,output] = linprog(...)
[x,fval,exitflag,output,lambda] = linprog(...)
```

Description

linprog solves linear programming problems.

$x = \text{linprog}(f,A,b)$ solves $\min f' \cdot x$ such that $A \cdot x \leq b$.

$x = \text{linprog}(f,A,b,Aeq,beq)$ solves the problem above while additionally satisfying the equality constraints $Aeq \cdot x = beq$. Set $A = []$ and $b = []$ if no inequalities exist.

$x = \text{linprog}(f,A,b,Aeq,beq,lb,ub)$ defines a set of lower and upper bounds on the design variables, x , so that the solution is always in the range $lb \leq x \leq ub$. Set $Aeq = []$ and $beq = []$ if no equalities exist.

$x = \text{linprog}(f,A,b,Aeq,beq,lb,ub,x0)$ sets the starting point to $x0$. linprog uses $x0$ only with the active-set algorithm. linprog ignores $x0$ with the interior-point and simplex algorithms.

`x = linprog(f,A,b,Aeq,beq,lb,ub,x0,options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options.

`x = linprog(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-192.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,fval] = linprog(...)` returns the value of the objective function `fun` at the solution `x`: $fval = f' * x$.

`[x,fval,exitflag] = linprog(...)` returns a value `exitflag` that describes the exit condition.

`[x,fval,exitflag,output] = linprog(...)` returns a structure `output` that contains information about the optimization.

`[x,fval,exitflag,output,lambda] = linprog(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the output `fval` is `[]`.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `linprog`. “Options” on page 10-194 provides the function-specific details for the options values.

problem	f	Linear objective function vector f
	Aineq	Matrix for linear inequality constraints
	bineq	Vector for linear inequality constraints
	Aeq	Matrix for linear equality constraints
	beq	Vector for linear equality constraints
	lb	Vector of lower bounds
	ub	Vector of upper bounds
	x0	Initial point for x, active set algorithm only
	solver	'linprog'
	options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `linprog`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.
1	Function converged to a solution x.
0	Number of iterations exceeded <code>options.MaxIter</code> .
-2	No feasible point was found.
-3	Problem is unbounded.
-4	NaN value was encountered during execution of the algorithm.
-5	Both primal and dual problems are infeasible.

	-7	Search direction became too small. No further progress could be made.
lambda		Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields of the structure are: lower Lower bounds lb upper Upper bounds ub ineqlin Linear inequalities eqlin Linear equalities
output		Structure containing information about the optimization. The fields of the structure are: iterations Number of iterations algorithm Optimization algorithm used cgiterations 0 (interior-point algorithm only, included for backward compatibility) message Exit message constrviolation Maximum of constraint functions firstorderopt First-order optimality measure

Options

Optimization options used by `linprog`. Some options apply to all algorithms, and others are only relevant when using the interior-point algorithm. Use `optimoptions` to set or change options. See “Optimization Options Reference” on page 9-7 for detailed information.

All Algorithms

All `linprog` algorithms use the following options:

Algorithm	Choose the optimization algorithm: <ul style="list-style-type: none">• 'interior-point' (default)• 'active-set'• 'simplex' <p>For information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.</p>
Diagnostics	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default 'off'.
Display	Level of display. <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'iter' displays output at each iteration. The 'iter' option only works with the interior-point and simplex algorithms.• 'final' (default) displays just the final output.
LargeScale Use Algorithm instead	Use large-scale algorithm when set to 'on' (default). Use a medium-scale algorithm when set to 'off' (see Simplex in “simplex Algorithm Only” on page 10-196). For information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.

MaxIter	Maximum number of iterations allowed, a positive integer. The default is: <ul style="list-style-type: none">• 85 for the interior-point algorithm• $10 \times \text{numberOfVariables}$ for the simplex algorithm• $10 \times \max(\text{numberOfVariables}, \text{numberOfInequalities} + \text{numberOfBounds})$ for the active-set algorithm
TolFun	Termination tolerance on the function value, a positive scalar. The default is: <ul style="list-style-type: none">• $1e-8$ for the interior-point algorithm• $1e-6$ for the simplex algorithm• The option is not used for the medium-scale active-set algorithm

simplex Algorithm Only

The simplex algorithm use the following option:

Simplex Use Algorithm instead	If 'on', and if LargeScale is 'off', linprog uses the simplex algorithm. The simplex algorithm uses a built-in starting point, ignoring the starting point x0 if supplied. The default is 'off', meaning linprog uses an active-set algorithm. See “Active-Set and Simplex Algorithms” on page 10-198 for more information and an example.
-------------------------------------	--

Examples

Find x that minimizes

$$f(x) = -5x_1 - 4x_2 - 6x_3,$$

subject to

$$\begin{aligned}x_1 - x_2 + x_3 &\leq 20 \\ 3x_1 + 2x_2 + 4x_3 &\leq 42 \\ 3x_1 + 2x_2 &\leq 30 \\ 0 \leq x_1, 0 \leq x_2, 0 \leq x_3.\end{aligned}$$

First, enter the coefficients

```
f = [-5; -4; -6];
A = [1 -1 1
     3 2 4
     3 2 0];
b = [20; 42; 30];
lb = zeros(3,1);
```

Next, call a linear programming routine.

```
[x,fval,exitflag,output,lambda] = linprog(f,A,b,[],[],lb);
```

Examine the solution and Lagrange multipliers:

```
x,lambda.ineqlin,lambda.lower
```

```
x =
    0.0000
   15.0000
    3.0000
```

```
ans =
    0.0000
    1.5000
    0.5000
```

```
ans =
    1.0000
    0.0000
    0.0000
```

Nonzero elements of the vectors in the fields of `lambda` indicate active constraints at the solution. In this case, the second and third inequality constraints (in `lambda.ineqlin`) and the first lower bound constraint (in `lambda.lower`) are active constraints (i.e., the solution is on their constraint boundaries).

Algorithms **Interior-Point Algorithm**

The interior-point method is based on LIPSOL (Linear Interior Point Solver, [3]), which is a variant of Mehrotra's predictor-corrector algorithm ([2]), a primal-dual interior-point method. A number of preprocessing steps occur before the algorithm begins to iterate. See "Interior-Point Linear Programming" on page 6-106.

Active-Set and Simplex Algorithms

`linprog` uses a projection method as used in the `quadprog` algorithm. `linprog` is an active set method and is thus a variation of the well-known *simplex* method for linear programming [1]. The algorithm finds an initial feasible solution by first solving another linear programming problem.

Alternatively, you can use the simplex algorithm, described in "linprog Simplex Algorithm" on page 6-114, by entering

```
options = optimoptions('linprog','Algorithm','simplex')
```

and passing `options` as an input argument to `linprog`. The simplex algorithm returns a vertex optimal solution.

Note `linprog` ignores `x0`, and computes its own initial point for the interior-point and simplex algorithms. `linprog` uses `x0` only with the active-set algorithm.

Diagnostics **Interior-Point Algorithm**

The first stage of the algorithm might involve some preprocessing of the constraints (see "Interior-Point Linear Programming" on page 6-106).

Several possible conditions might occur that cause `linprog` to exit with an infeasibility message. In each case, the `exitflag` argument returned by `linprog` is set to a negative value to indicate failure.

If a row of all zeros is detected in `Aeq` but the corresponding element of `beq` is not zero, the exit message is

Exiting due to infeasibility: An all-zero row in the constraint matrix does not have a zero in corresponding right-hand-side entry.

If one of the elements of `x` is found not to be bounded below, the exit message is

Exiting due to infeasibility: Objective $f'*x$ is unbounded below.

If one of the rows of `Aeq` has only one nonzero element, the associated value in `x` is called a *singleton* variable. In this case, the value of that component of `x` can be computed from `Aeq` and `beq`. If the value computed violates another constraint, the exit message is

Exiting due to infeasibility: Singleton variables in equality constraints are not feasible.

If the singleton variable can be solved for but the solution violates the upper or lower bounds, the exit message is

Exiting due to infeasibility: Singleton variables in the equality constraints are not within bounds.

Note The preprocessing steps are cumulative. For example, even if your constraint matrix does not have a row of all zeros to begin with, other preprocessing steps may cause such a row to occur.

Once the preprocessing has finished, the iterative part of the algorithm begins until the stopping criteria are met. (See “Interior-Point Linear

Programming” on page 6-106 for more information about residuals, the primal problem, the dual problem, and the related stopping criteria.) If the residuals are growing instead of getting smaller, or the residuals are neither growing nor shrinking, one of the two following termination messages is displayed, respectively,

One or more of the residuals, duality gap, or total relative error has grown 100000 times greater than its minimum value so far:

OR

One or more of the residuals, duality gap, or total relative error has stalled:

After one of these messages is displayed, it is followed by one of the following six messages indicating that the dual, the primal, or both appear to be infeasible. The messages differ according to how the infeasibility or unboundedness was measured.

The dual appears to be infeasible (and the primal unbounded).(The primal residual < TolFun.)

The primal appears to be infeasible (and the dual unbounded). (The dual residual < TolFun.)

The dual appears to be infeasible (and the primal unbounded) since the dual residual > sqrt(TolFun).(The primal residual < 10*TolFun.)

The primal appears to be infeasible (and the dual unbounded) since the primal residual > sqrt(TolFun).(The dual residual < 10*TolFun.)

The dual appears to be infeasible and the primal unbounded since the primal objective < -1e+10 and the dual objective < 1e+6.

The primal appears to be infeasible and the dual unbounded since the dual objective > 1e+10 and the primal objective > -1e+6.

Both the primal and the dual appear to be infeasible.

Note that, for example, the primal (objective) can be unbounded and the primal residual, which is a measure of primal constraint satisfaction, can be small.

Active-Set and Simplex Algorithms

linprog gives a warning when the problem is infeasible.

Warning: The constraints are overly stringent;
there is no feasible solution.

In this case, linprog produces a result that minimizes the worst case constraint violation.

When the equality constraints are inconsistent, linprog gives

Warning: The equality constraints are overly
stringent; there is no feasible solution.

Unbounded solutions result in the warning

Warning: The solution is unbounded and at infinity;
the constraints are not restrictive enough.

In this case, linprog returns a value of x that satisfies the constraints.

Limitations

Active-Set Algorithm

At this time, the only levels of display, using the `Display` option in options, are 'off' and 'final'; iterative output using 'iter' is not available.

Interior-Point Algorithm

Coverage and Requirements

For Large Problems
A and Aeq should be sparse.

References

[1] Dantzig, G.B., A. Orden, and P. Wolfe, "Generalized Simplex Method for Minimizing a Linear Form Under Linear Inequality Restraints," *Pacific Journal Math.*, Vol. 5, pp. 183–195, 1955.

[2] Mehrotra, S., “On the Implementation of a Primal-Dual Interior Point Method,” *SIAM Journal on Optimization*, Vol. 2, pp. 575–601, 1992.

[3] Zhang, Y., “Solving Large-Scale Linear Programs by Interior-Point Methods Under the MATLAB Environment,” *Technical Report TR96-01*, Department of Mathematics and Statistics, University of Maryland, Baltimore County, Baltimore, MD, July 1995.

See Also

`intlinprog` | `quadprog` | `optimtool`

How To

- “Linear Programming and Mixed-Integer Linear Programming”

Purpose Solve nonlinear curve-fitting (data-fitting) problems in least-squares sense

Equation Find coefficients x that solve the problem

$$\min_x \|F(x, xdata) - ydata\|_2^2 = \min_x \sum_i (F(x, xdata_i) - ydata_i)^2,$$

given input data $xdata$, and the observed output $ydata$, where $xdata$ and $ydata$ are matrices or vectors, and $F(x, xdata)$ is a matrix-valued or vector-valued function of the same size as $ydata$.

Optionally, the components of x can have lower and upper bounds lb , and ub . x , lb , and ub can be vectors or matrices; see “Matrix Arguments” on page 2-32.

The `lsqcurvefit` function uses the same algorithm as `lsqnonlin`. `lsqcurvefit` simply provides a convenient interface for data-fitting problems.

Syntax

```
x = lsqcurvefit(fun,x0,xdata,ydata)
x = lsqcurvefit(fun,x0,xdata,ydata,lb,ub)
x = lsqcurvefit(fun,x0,xdata,ydata,lb,ub,options)
x = lsqcurvefit(problem)
[x,resnorm] = lsqcurvefit(...)
[x,resnorm,residual] = lsqcurvefit(...)
[x,resnorm,residual,exitflag] = lsqcurvefit(...)
[x,resnorm,residual,exitflag,output] = lsqcurvefit(...)
[x,resnorm,residual,exitflag,output,lambda] = lsqcurvefit(...)
[x,resnorm,residual,exitflag,output,lambda,
 jacobian] = lsqcurvefit(...)
```

Description

`x = lsqcurvefit(fun,x0,xdata,ydata)` starts at x_0 and finds coefficients x to best fit the nonlinear function $\text{fun}(x, xdata)$ to the data $ydata$ (in the least-squares sense). $ydata$ must be the same size as the vector (or matrix) F returned by `fun`.

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to `fun`, if necessary.

`fun` should return `fun(x,xdata)`, and not the sum-of-squares `sum((fun(x,xdata)-ydata).^2)`. `lsqcurvefit` implicitly computes the sum of squares of the components of `fun(x,xdata)-ydata`.

`x = lsqcurvefit(fun,x0,xdata,ydata,lb,ub)` defines a set of lower and upper bounds on the design variables in `x` so that the solution is always in the range $lb \leq x \leq ub$.

`x = lsqcurvefit(fun,x0,xdata,ydata,lb,ub,options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options. Pass empty matrices for `lb` and `ub` if no bounds exist.

`x = lsqcurvefit(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-205.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,resnorm] = lsqcurvefit(...)` returns the value of the squared 2-norm of the residual at `x`: `sum((fun(x,xdata)-ydata).^2)`.

`[x,resnorm,residual] = lsqcurvefit(...)` returns the value of the residual `fun(x,xdata)-ydata` at the solution `x`.

`[x,resnorm,residual,exitflag] = lsqcurvefit(...)` returns a value `exitflag` that describes the exit condition.

`[x,resnorm,residual,exitflag,output] = lsqcurvefit(...)` returns a structure `output` that contains information about the optimization.

`[x,resnorm,residual,exitflag,output,lambda] = lsqcurvefit(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

`[x,resnorm,residual,exitflag,output,lambda,jacobian] = lsqcurvefit(...)` returns the Jacobian of `fun` at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the outputs `resnorm` and `residual` are `[]`.

Components of `x0` that violate the bounds `lb < x < ub` are reset to the interior of the box defined by the bounds. Components that respect the bounds are not changed.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `lsqcurvefit`. This section provides function-specific details for `fun`, options, and problem:

`fun` The function you want to fit. `fun` is a function that takes two inputs: a vector or matrix `x`, and a vector or matrix `xdata`. `fun` returns a vector or matrix `F`, the objective function evaluated at `x` and `xdata`. The function `fun` can be specified as a function handle for a function file:

```
x = lsqcurvefit(@myfun,x0,xdata,ydata)
```

where `myfun` is a MATLAB function such as

```
function F = myfun(x,xdata)
F = ...        % Compute function values at x, xdata
```

`fun` can also be a function handle for an anonymous function.

```
f = @(x,xdata)x(1)*xdata.^2+x(2)*sin(xdata);
x = lsqcurvefit(f,x0,xdata,ydata);
```

`lsqcurvefit` internally converts matrix `x` or `F` to vectors using linear indexing.

Note fun should return `fun(x,xdata)`, and not the sum-of-squares `sum((fun(x,xdata)-ydata).^2)`. `lsqcurvefit` implicitly computes the sum of squares of the components of `fun(x,xdata)-ydata`.

If the Jacobian can also be computed *and* the Jacobian option is 'on', set by

```
options = optimoptions('lsqcurvefit','Jacobian','on')
```

then the function `fun` must return, in a second output argument, the Jacobian value `J`, a matrix, at `x`. By checking the value of `nargout`, the function can avoid computing `J` when `fun` is called with only one output argument (in the case where the optimization algorithm only needs the value of `F` but not `J`).

```
function [F,J] = myfun(x,xdata)
F = ...           % objective function values at x
if nargout > 1   % two output arguments
    J = ...       % Jacobian of the function evaluated at x
end
```

If `fun` returns a vector (matrix) of `m` components and `x` has length `n`, where `n` is the length of `x0`, then the Jacobian `J` is an `m`-by-`n` matrix where `J(i,j)` is the partial derivative of `F(i)` with respect to `x(j)`. (The Jacobian `J` is the transpose of the gradient of `F`.) For more information, see “Writing Vector and Matrix Objective Functions” on page 2-27.

`options` “Options” on page 10-209 provides the function-specific details for the options values.

problem	objective	Objective function of x and xdata
	x0	Initial point for x, active set algorithm only
	xdata	Input data for objective function
	ydata	Output data to be matched by objective function
	lb	Vector of lower bounds
	ub	Vector of upper bounds
	solver	'lsqcurvefit'
	options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `lsqcurvefit`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated:
1	Function converged to a solution <code>x</code> .
2	Change in <code>x</code> was less than the specified tolerance.
3	Change in the residual was less than the specified tolerance.
4	Magnitude of search direction smaller than the specified tolerance.
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .

	-1	Output function terminated the algorithm.
	-2	Problem is infeasible: the bounds lb and ub are inconsistent.
	-4	Optimization could not make further progress.
lambda		Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields of the structure are
	lower	Lower bounds lb
	upper	Upper bounds ub
output		Structure containing information about the optimization. The fields of the structure are
	firstorderopt	Measure of first-order optimality (trust-region-reflective algorithm, [] for others).
	iterations	Number of iterations taken
	funcCount	Number of function evaluations
	cgiterations	Total number of PCG iterations (trust-region-reflective algorithm, [] for others)
	algorithm	Optimization algorithm used
	stepsize	Final displacement in x (Levenberg-Marquardt algorithm).
	message	Exit message

Note The sum of squares should not be formed explicitly. Instead, your function should return a vector of function values. See the examples below.

Options

Optimization options used by `lsqcurvefit`. Some options apply to all algorithms, some are only relevant when using the trust-region-reflective algorithm, and others are only relevant when you are using the Levenberg-Marquardt algorithm. Use `optimoptions` to set or change options. See “Algorithm Options” on page 10-209 for detailed information.

The `Algorithm` option specifies a preference for which algorithm to use. It is only a preference, because certain conditions must be met to use the trust-region-reflective or Levenberg-Marquardt algorithm. For the trust-region-reflective algorithm, the nonlinear system of equations cannot be underdetermined; that is, the number of equations (the number of elements of `F` returned by `fun`) must be at least as many as the length of `x`. Furthermore, only the trust-region-reflective algorithm handles bound constraints.

Algorithm Options

Both algorithms use the following option:

Algorithm	Choose between 'trust-region-reflective' (default) and 'levenberg-marquardt'. Set the initial Levenberg-Marquardt parameter λ by setting <code>Algorithm</code> to a cell array such as <code>{'levenberg-marquardt', .005}</code> . The default $\lambda = 0.01$.
-----------	---

The `Algorithm` option specifies a preference for which algorithm to use. It is only a preference, because certain conditions must be met to use each algorithm. For the trust-region-reflective algorithm, the nonlinear system of equations cannot be underdetermined; that is, the number of

equations (the number of elements of F returned by `fun`) must be at least as many as the length of x . The Levenberg-Marquardt algorithm does not handle bound constraints. For more information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.

<code>DerivativeCheck</code>	Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives. The choices are 'on' or the default 'off'.
<code>Diagnostics</code>	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default 'off'.
<code>DiffMaxChange</code>	Maximum change in variables for finite-difference gradients (a positive scalar). The default is <code>Inf</code> .
<code>DiffMinChange</code>	Minimum change in variables for finite-difference gradients (a positive scalar). The default is <code>0</code> .
<code>Display</code>	Level of display: <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'iter' displays output at each iteration, and gives the default exit message.• 'iter-detailed' displays output at each iteration, and gives the technical exit message.• 'final' (default) displays just the final output, and gives the default exit message.• 'final-detailed' displays just the final output, and gives the technical exit message.

FinDiffRelStep Scalar or vector step size factor. When you set `FinDiffRelStep` to a vector `v`, forward finite differences `delta` are

```
delta =
v.*sign(x).*max(abs(x),TypicalX);
```

and central finite differences are

```
delta = v.*max(abs(x),TypicalX);
```

Scalar `FinDiffRelStep` expands to a vector. The default is `sqrt(eps)` for forward finite differences, and `eps^(1/3)` for central finite differences.

FinDiffType Finite differences, used to estimate gradients, are either 'forward' (default), or 'central' (centered). 'central' takes twice as many function evaluations, but should be more accurate.

The algorithm is careful to obey bounds when estimating both types of finite differences. So, for example, it could take a backward, rather than a forward, difference to avoid evaluating at a point outside bounds.

FunValCheck Check whether function values are valid. 'on' displays an error when the function returns a value that is complex, Inf, or NaN. The default 'off' displays no error.

Jacobian If 'on', `lsqcurvefit` uses a user-defined Jacobian (defined in `fun`), or Jacobian information (when using `JacobMult`), for the objective function. If 'off' (default), `lsqcurvefit` approximates the Jacobian using finite differences.

MaxFunEvals	Maximum number of function evaluations allowed, a positive integer. The default is $100 \times \text{numberOfVariables}$.
MaxIter	Maximum number of iterations allowed, a positive integer. The default is 400.
OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none (<code>[]</code>). See “Output Function” on page 9-21.
PlotFcns	Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none (<code>[]</code>): <ul style="list-style-type: none">• <code>@optimplotx</code> plots the current point.• <code>@optimplotfunccount</code> plots the function count.• <code>@optimplotfval</code> plots the function value.• <code>@optimplotresnorm</code> plots the norm of the residuals.• <code>@optimplotstepsize</code> plots the step size.• <code>@optimplotfirstorderopt</code> plots the first-order optimality measure. For information on writing a custom plot function, see “Plot Functions” on page 9-30.
TolFun	Termination tolerance on the function value, a positive scalar. The default is $1e-6$.

<code>TolX</code>	Termination tolerance on x , a positive scalar. The default is $1e-6$.
<code>TypicalX</code>	Typical x values. The number of elements in <code>TypicalX</code> is equal to the number of elements in <code>x0</code> , the starting point. The default value is <code>ones(numberofvariables,1)</code> . <code>lsqcurvefit</code> uses <code>TypicalX</code> for scaling finite differences for gradient estimation.

Trust-Region-Reflective Algorithm Only

The trust-region-reflective algorithm uses the following options:

<code>JacobMult</code>	Function handle for Jacobian multiply function. For large-scale structured problems, this function computes the Jacobian matrix product $J*Y$, $J' * Y$, or $J' * (J*Y)$ without actually forming J . The function is of the form
------------------------	---

$$W = \text{jmfun}(J\text{info}, Y, \text{flag})$$

where `Jinfo` contains the matrix used to compute $J*Y$ (or $J' * Y$, or $J' * (J*Y)$). The first argument `Jinfo` must be the same as the second argument returned by the objective function `fun`, for example, in

$$[F, J\text{info}] = \text{fun}(x)$$

Y is a matrix that has the same number of rows as there are dimensions in the problem. `flag` determines which product to compute:

- If `flag == 0` then $W = J' * (J*Y)$.
- If `flag > 0` then $W = J*Y$.
- If `flag < 0` then $W = J' * Y$.

In each case, J is not formed explicitly. `lsqcurvefit` uses `Jinfo` to compute the preconditioner. See “Passing Extra Parameters” on page 2-53 for information on how to supply values for any additional parameters `jmfun` needs.

Note 'Jacobian' must be set to 'on' for `lsqcurvefit` to pass `Jinfo` from `fun` to `jmfun`.

See “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80 and “Jacobian Multiply Function with Linear Least Squares” on page 6-224 for similar examples.

JacobPattern

Sparsity pattern of the Jacobian for finite differencing. Set `JacobPattern(i,j) = 1` when `fun(i)` depends on `x(j)`. Otherwise, set `JacobPattern(i,j) = 0`. In other words, `JacobPattern(i,j) = 1` when you can have $\partial \text{fun}(i) / \partial x(j) \neq 0$.

Use `JacobPattern` when it is inconvenient to compute the Jacobian matrix J in `fun`, though you can determine (say, by inspection) when `fun(i)` depends on `x(j)`. `lsqcurvefit` can approximate J via sparse finite differences when you give `JacobPattern`.

In the worst case, if the structure is unknown, do not set `JacobPattern`. The default behavior is as if `JacobPattern` is a dense matrix of ones. Then `lsqcurvefit` computes a full finite-difference approximation in each iteration. This can be

	very expensive for large problems, so it is usually better to determine the sparsity structure.
MaxPCGIter	Maximum number of PCG (preconditioned conjugate gradient) iterations, a positive scalar. The default is <code>max(1, floor(numberOfVariables/2))</code> . For more information, see “Algorithms” on page 10-217.
PrecondBandWidth	Upper bandwidth of preconditioner for PCG, a nonnegative integer. The default PrecondBandWidth is Inf, which means a direct factorization (Cholesky) is used rather than the conjugate gradients (CG). The direct factorization is computationally more expensive than CG, but produces a better quality step towards the solution. Set PrecondBandWidth to 0 for diagonal preconditioning (upper bandwidth of 0). For some problems, an intermediate bandwidth reduces the number of PCG iterations.
TolPCG	Termination tolerance on the PCG iteration, a positive scalar. The default is 0.1.

Levenberg-Marquardt Algorithm Only

The Levenberg-Marquardt algorithm uses the following options:

ScaleProblem	'Jacobian' can sometimes improve the convergence of a poorly-scaled problem; the default is 'none'.
--------------	---

Examples

Given vectors of data *xdata* and *ydata*, suppose you want to find coefficients *x* to find the best fit to the exponential decay equation

$$ydata(i) = x(1)e^{x(2)xdata(i)}$$

That is, you want to minimize

$$\min_x \sum_i (F(x, xdata_i) - ydata_i)^2,$$

where m is the length of $xdata$ and $ydata$, the function F is defined by

```
F(x, xdata) = x(1)*exp(x(2)*xdata);
```

and the starting point is $x0 = [100; -1];$.

First, write a file to return the value of F (F has n components).

```
function F = myfun(x, xdata)
F = x(1)*exp(x(2)*xdata);
```

Next, invoke an optimization routine:

```
% Assume you determined xdata and ydata experimentally
xdata = ...
[0.9 1.5 13.8 19.8 24.1 28.2 35.2 60.3 74.6 81.3];
ydata = ...
[455.2 428.6 124.1 67.3 43.2 28.1 13.1 -0.4 -1.3 -1.5];
x0 = [100; -1] % Starting guess
[x, resnorm] = lsqcurvefit(@myfun, x0, xdata, ydata);
```

At the time that `lsqcurvefit` is called, $xdata$ and $ydata$ are assumed to exist and are vectors of the same size. They must be the same size because the value F returned by `fun` must be the same size as $ydata$.

After 27 function evaluations, this example gives the solution

```
x, resnorm

x =
  498.8309 -0.1013
resnorm =
  9.5049
```

There may be a slight variation in the number of iterations and the value of the returned x , depending on the platform and release.

Algorithms

Trust-Region-Reflective Optimization

By default `lsqcurvefit` chooses the trust-region-reflective algorithm. This algorithm is a subspace trust-region method and is based on the interior-reflective Newton method described in [1] and [2]. Each iteration involves the approximate solution of a large linear system using the method of preconditioned conjugate gradients (PCG). See “Trust-Region-Reflective Least Squares” on page 6-201, and in particular, “Large Scale Nonlinear Least Squares” on page 6-203.

Levenberg-Marquardt Optimization

If you set the `Algorithm` option to 'levenberg-marquardt' with `optimoptions`, `lsqcurvefit` uses the Levenberg-Marquardt method [4], [5], and [6]. See “Levenberg-Marquardt Method” on page 6-205.

Diagnostics

Trust-Region-Reflective Optimization

The trust-region-reflective method does not allow equal upper and lower bounds. For example, if `lb(2)==ub(2)`, `lsqcurvefit` gives the error

```
Equal upper and lower bounds not permitted.
```

`lsqcurvefit` does not handle equality constraints, which is another way to formulate equal bounds. If equality constraints are present, use `fmincon`, `fminimax`, or `fgoalattain` for alternative formulations where equality constraints can be included.

Limitations

The function to be minimized must be continuous. `lsqcurvefit` might only give local solutions.

`lsqcurvefit` can solve complex-valued problems directly with the `levenberg-marquardt` algorithm. However, this algorithm does not accept bound constraints. For a complex problem with bound constraints, split the variables into real and imaginary parts, and

use the trust-region-reflective algorithm. See “Fit a Model to Complex-Valued Data” on page 6-231.

Note The Statistics Toolbox function `nlinfit` has more statistics-oriented outputs that are useful, for example, in finding confidence intervals for the coefficients. It also comes with the `nlintool` GUI for visualizing the fitted function.

The `lsqnonlin` function has more outputs related to how well the optimization performed. It can put bounds on the parameters, and it accepts many options to control the optimization algorithm.

Trust-Region-Reflective Optimization

The trust-region-reflective algorithm for `lsqcurvefit` does not solve underdetermined systems; it requires that the number of equations, i.e., the row dimension of F , be at least as great as the number of variables. In the underdetermined case, the Levenberg-Marquardt algorithm is used instead.

The preconditioner computation used in the preconditioned conjugate gradient part of the trust-region-reflective method forms $J^T J$ (where J is the Jacobian matrix) before computing the preconditioner; therefore, a row of J with many nonzeros, which results in a nearly dense product $J^T J$, can lead to a costly solution process for large problems.

If components of x have no upper (or lower) bounds, then `lsqcurvefit` prefers that the corresponding components of `ub` (or `lb`) be set to `inf` (or `-inf` for lower bounds) as opposed to an arbitrary but very large positive (or negative for lower bounds) number.

Trust-Region-Reflective Problem Coverage and Requirements

For Large Problems

- Provide sparsity structure of the Jacobian or compute the Jacobian in fun.
- The Jacobian should be sparse.

Levenberg-Marquardt Optimization

The Levenberg-Marquardt algorithm does not handle bound constraints.

Since the trust-region-reflective algorithm does not handle underdetermined systems and the Levenberg-Marquardt does not handle bound constraints, problems with both these characteristics cannot be solved by `lsqcurvefit`.

References

- [1] Coleman, T.F. and Y. Li, "An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds," *SIAM Journal on Optimization*, Vol. 6, pp. 418-445, 1996.
- [2] Coleman, T.F. and Y. Li, "On the Convergence of Reflective Newton Methods for Large-Scale Nonlinear Minimization Subject to Bounds," *Mathematical Programming*, Vol. 67, Number 2, pp. 189-224, 1994.
- [3] Dennis, J. E. Jr., "Nonlinear Least-Squares," *State of the Art in Numerical Analysis*, ed. D. Jacobs, Academic Press, pp. 269-312, 1977.
- [4] Levenberg, K., "A Method for the Solution of Certain Problems in Least-Squares," *Quarterly Applied Math.* 2, pp. 164-168, 1944.
- [5] Marquardt, D., "An Algorithm for Least-Squares Estimation of Nonlinear Parameters," *SIAM Journal Applied Math.*, Vol. 11, pp. 431-441, 1963.
- [6] More, J. J., "The Levenberg-Marquardt Algorithm: Implementation and Theory," *Numerical Analysis*, ed. G. A. Watson, Lecture Notes in Mathematics 630, Springer Verlag, pp. 105-116, 1977.

lsqcurvefit

See Also

\ | lsqlin | lsqnonlin | lsqnonneg | optimoptions | optimtool
| nlinfit

How To

- `function_handle`
- “Least-Squares (Model Fitting) Algorithms” on page 6-200
- “Nonlinear Curve Fitting with `lsqcurvefit`” on page 6-229

Purpose Solve constrained linear least-squares problems

Equation Solves least-squares curve fitting problems of the form

$$\min_x \frac{1}{2} \|C \cdot x - d\|_2^2 \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ lb \leq x \leq ub. \end{cases}$$

Syntax

```
x = lsqlin(C,d,A,b)
x = lsqlin(C,d,A,b,Aeq,beq)
x = lsqlin(C,d,A,b,Aeq,beq,lb,ub)
x = lsqlin(C,d,A,b,Aeq,beq,lb,ub,x0)
x = lsqlin(C,d,A,b,Aeq,beq,lb,ub,x0,options)
x = lsqlin(problem)
[x,resnorm] = lsqlin(...)
[x,resnorm,residual] = lsqlin(...)
[x,resnorm,residual,exitflag] = lsqlin(...)
[x,resnorm,residual,exitflag,output] = lsqlin(...)
[x,resnorm,residual,exitflag,output,lambda] = lsqlin(...)
```

Description

`x = lsqlin(C,d,A,b)` solves the linear system $C \cdot x = d$ in the least-squares sense subject to $A \cdot x \leq b$, where C is m-by-n.

`x = lsqlin(C,d,A,b,Aeq,beq)` solves the preceding problem while additionally satisfying the equality constraints $Aeq \cdot x = beq$. Set $A = []$ and $b = []$ if no inequalities exist.

`x = lsqlin(C,d,A,b,Aeq,beq,lb,ub)` defines a set of lower and upper bounds on the design variables in x so that the solution is always in the range $lb \leq x \leq ub$. Set $Aeq = []$ and $beq = []$ if no equalities exist.

`x = lsqlin(C,d,A,b,Aeq,beq,lb,ub,x0)` sets the starting point to $x0$. Set $lb = []$ and $ub = []$ if no bounds exist.

`x = lsqlin(C,d,A,b,Aeq,beq,lb,ub,x0,options)` minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options.

`x = lsqlin(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-223.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,resnorm] = lsqlin(...)` returns the value of the squared 2-norm of the residual, $\text{norm}(C*x-d)^2$.

`[x,resnorm,residual] = lsqlin(...)` returns the residual $C*x-d$.

`[x,resnorm,residual,exitflag] = lsqlin(...)` returns a value `exitflag` that describes the exit condition.

`[x,resnorm,residual,exitflag,output] = lsqlin(...)` returns a structure `output` that contains information about the optimization.

`[x,resnorm,residual,exitflag,output,lambda] = lsqlin(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output x is x_0 and the outputs `resnorm` and `residual` are `[]`.

Components of x_0 that violate the bounds $lb \leq x \leq ub$ are reset to the interior of the box defined by the bounds. Components that respect the bounds are not changed.

If no x_0 is provided, x_0 is set to the zero vector. If any component of this zero vector x_0 violates the bounds, x_0 is set to a point in the interior of the box defined by the bounds.

The factor $\frac{1}{2}$ in the definition of the problem affects the values in the `lambda` structure.

You can solve some large structured problems, including those where the `C` matrix is too large to fit in memory, using the trust-region-reflective algorithm with a Jacobian multiply function. For information, see “Trust-Region-Reflective Algorithm Only” on page 10-226.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `lsqlin`. “Options” on page 10-225 provides the options values specific to `lsqlin`.

<code>problem</code>	<code>C</code>	Matrix
	<code>d</code>	Vector
	<code>Aineq</code>	Matrix for linear inequality constraints
	<code>bineq</code>	Vector for linear inequality constraints
	<code>Aeq</code>	Matrix for linear equality constraints
	<code>beq</code>	Vector for linear equality constraints
	<code>lb</code>	Vector of lower bounds
	<code>ub</code>	Vector of upper bounds

x0	Initial point for x
solver	'lsqlin'
options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `lsqlin`. This section provides function-specific details for `exitflag`, `lambda`, and output:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated:
1	Function converged to a solution <code>x</code> .
3	Change in the residual was smaller than the specified tolerance.
0	Number of iterations exceeded <code>options.MaxIter</code> .
-2	The problem is infeasible.
-4	Ill-conditioning prevents further optimization.
-7	Magnitude of search direction became too small. No further progress could be made.

<code>lambda</code>	Structure containing the Lagrange multipliers at the solution <code>x</code> (separated by constraint type). The fields are
<code>lower</code>	Lower bounds <code>lb</code>
<code>upper</code>	Upper bounds <code>ub</code>
<code>ineqlin</code>	Linear inequalities
<code>eqlin</code>	Linear equalities
<code>output</code>	Structure containing information about the optimization. The fields are
<code>iterations</code>	Number of iterations taken
<code>algorithm</code>	Optimization algorithm used
<code>cgiterations</code>	Total number of PCG iterations (trust-region-reflective algorithm, [] for active-set)
<code>firstorderopt</code>	Measure of first-order optimality (trust-region-reflective algorithm, [] for active-set)
<code>message</code>	Exit message

Options

Optimization options used by `lsqlin`. Set or change the values of these options using the `optimoptions` function. Some options apply to all algorithms, some are only relevant when you are using the `trust-region-reflective` algorithm. See “Optimization Options Reference” on page 9-7 for detailed information.

All Algorithms

All algorithms use the following options:

Algorithm	Choose the lsqlin algorithm. Choices are 'active-set' and 'trust-region-reflective' (default).
If you use optimset (not recommended), use LargeScale instead of Algorithm.	The trust-region-reflective algorithm requires <i>only</i> upper and lower bounds, meaning no linear inequalities or equalities. Otherwise, lsqlin uses the active-set algorithm. For more information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.
Diagnostics	Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default 'off'.
Display	Level of display. 'off' or 'none' displays no output; 'final' (default) displays just the final output.
LargeScale	Use the large-scale algorithm if possible when set to 'on' (default). Use the medium-scale algorithm when set to 'off'.
If you use optimoptions (recommended), use Algorithm instead of LargeScale.	The large-scale algorithm requires <i>only</i> upper and lower bounds, meaning no linear inequalities or equalities. Otherwise, lsqlin uses the medium-scale algorithm. For more information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.
MaxIter	Maximum number of iterations allowed, a positive integer. The default value is 200.

Trust-Region-Reflective Algorithm Only

The trust-region-reflective algorithm uses the following options:

JacobMult

Function handle for Jacobian multiply function. For large-scale structured problems, this function should compute the Jacobian matrix product $C*Y$, $C'*Y$, or $C'*(C*Y)$ without actually forming C . Write the function in the form

$W = \text{jmfun}(\text{Jinfo}, Y, \text{flag})$

where Jinfo contains a matrix used to compute $C*Y$ (or $C'*Y$, or $C'*(C*Y)$).

jmfun must compute one of three different products, depending on the value of flag that lsqin passes:

- If $\text{flag} == 0$ then $W = C'*(C*Y)$.
- If $\text{flag} > 0$ then $W = C*Y$.
- If $\text{flag} < 0$ then $W = C'*Y$.

In each case, jmfun need not form C explicitly. lsqin uses Jinfo to compute the preconditioner. See “Passing Extra Parameters” on page 2-53 for information on how to supply extra parameters if necessary.

See “Jacobian Multiply Function with Linear Least Squares” on page 6-224 for an example.

MaxPCGIter

Maximum number of PCG (preconditioned conjugate gradient) iterations, a positive scalar. The default is $\max(1, \text{floor}(\text{numberOfVariables}/2))$. For more information, see “Algorithms” on page 10-231.

PrecondBandWidth	Upper bandwidth of preconditioner for PCG. By default, diagonal preconditioning is used (upper bandwidth of 0). For some problems, increasing the bandwidth reduces the number of PCG iterations. Setting PrecondBandWidth to Inf uses a direct factorization (Cholesky) rather than the conjugate gradients (CG). The direct factorization is computationally more expensive than CG, but produces a better quality step towards the solution.
TolFun	Termination tolerance on the function value, a positive scalar. The default is $100 \cdot \text{eps}$, about $2.2204e-14$.
TolPCG	Termination tolerance on the PCG iteration, a positive scalar. The default is 0.1.
TypicalX	Typical x values. The number of elements in TypicalX is equal to the number of elements in x0, the starting point. The default value is ones(numberofvariables,1). lsqlin uses TypicalX internally for scaling. TypicalX has an effect only when x has unbounded components, and when a TypicalX value for an unbounded component is larger than 1.

Examples

Find the least-squares solution to the overdetermined system $Cx = d$, subject to $Ax \leq b$ and $lb \leq x \leq ub$.

First, enter the coefficient matrices and the lower and upper bounds.

```
C = [  
    0.9501    0.7620    0.6153    0.4057  
    0.2311    0.4564    0.7919    0.9354  
    0.6068    0.0185    0.9218    0.9169  
    0.4859    0.8214    0.7382    0.4102  
    0.8912    0.4447    0.1762    0.8936];
```

```
d = [  
    0.0578  
    0.3528  
    0.8131  
    0.0098  
    0.1388];  
A = [  
    0.2027    0.2721    0.7467    0.4659  
    0.1987    0.1988    0.4450    0.4186  
    0.6037    0.0152    0.9318    0.8462];  
b = [  
    0.5251  
    0.2026  
    0.6721];  
lb = -0.1*ones(4,1);  
ub = 2*ones(4,1);
```

Next, call the constrained linear least-squares routine.

```
[x,resnorm,residual,exitflag,output,lambda] = ...  
    lsqlin(C,d,A,b,[ ],[ ],lb,ub);
```

Examine `x`, `lambda.ineqlin`, `lambda.lower`, and `lambda.upper`:

```
x  
x =  
   -0.1000  
   -0.1000  
    0.2152  
    0.3502  
  
lambda.ineqlin  
ans =  
     0  
    0.2392  
     0  
  
lambda.lower
```

```
ans =  
    0.0409  
    0.2784  
         0  
         0  
  
lambda.upper  
ans =  
         0  
         0  
         0  
         0
```

Nonzero elements of the vectors in the fields of `lambda` indicate active constraints at the solution. In this case, the second inequality constraint (in `lambda.ineqlin`) and the first lower and second lower bound constraints (in `lambda.lower`) are active constraints (i.e., the solution is on their constraint boundaries).

Notes

For problems with no constraints, use `\` (matrix left division). For example, `x = A\b`.

Because the problem being solved is always convex, `lsqlin` will find a global, although not necessarily unique, solution.

Better numerical results are likely if you specify equalities explicitly, using `Aeq` and `beq`, instead of implicitly, using `lb` and `ub`.

Trust-Region-Reflective Algorithm

If `x0` is not strictly feasible, `lsqlin` chooses a new strictly feasible (centered) starting point.

If components of x have no upper (or lower) bounds, set the corresponding components of `ub` (or `lb`) to `Inf` (or `-Inf` for `lb`) as opposed to an arbitrary but very large positive (or negative in the case of lower bounds) number.

Algorithms

Trust-Region-Reflective Algorithm

When the problem given to `lsqlin` has *only* upper and lower bounds; i.e., no linear inequalities or equalities are specified, and the matrix `C` has at least as many rows as columns, the default algorithm is `trust-region-reflective`. This method is a subspace trust-region method based on the interior-reflective Newton method described in [1]. Each iteration involves the approximate solution of a large linear system using the method of preconditioned conjugate gradients (PCG). See “Trust-Region Methods for Nonlinear Minimization” on page 6-5 and “Preconditioned Conjugate Gradient Method” on page 6-29.

Active-Set Algorithm

`lsqlin` uses the active-set algorithm when you specify it with `optimoptions`, or when you give linear inequalities or equalities. The algorithm is based on `quadprog`, which uses an active set method similar to that described in [2]. It finds an initial feasible solution by first solving a linear programming problem. See “trust-region-reflective `quadprog` Algorithm” on page 6-126.

Diagnostics

Trust-Region-Reflective Algorithm

The trust-region-reflective algorithm does not allow equal upper and lower bounds. For example, if `lb(2) == ub(2)`, then `lsqlin` gives the following error:

```
Equal upper and lower bounds not permitted  
in this large-scale method.  
Use equality constraints and the medium-scale  
method instead.
```

At this time, you must use the active-set algorithm to solve equality constrained problems.

Active-Set Algorithm

If the matrices `C`, `A`, or `Aeq` are sparse, and the problem formulation is not solvable using the trust-region-reflective algorithm, `lsqlin` warns that the matrices are converted to full.

lsqlin

Warning: This problem formulation not yet available for sparse matrices.
Converting to full to solve.

When a problem is infeasible, lsqlin gives a warning:

Warning: The constraints are overly stringent;
there is no feasible solution.

In this case, lsqlin produces a result that minimizes the worst case constraint violation.

When the equality constraints are inconsistent, lsqlin gives

Warning: The equality constraints are overly stringent;
there is no feasible solution.

Limitations

At this time, the only levels of display, using the Display option in options, are 'off' and 'final'; iterative output using 'iter' is not available.

Trust-Region-Reflective Algorithm Requirements

For Large Problems
C should be sparse.

References

[1] Coleman, T.F. and Y. Li, "A Reflective Newton Method for Minimizing a Quadratic Function Subject to Bounds on Some of the Variables," *SIAM Journal on Optimization*, Vol. 6, Number 4, pp. 1040-1058, 1996.

[2] Gill, P.E., W. Murray, and M.H. Wright, *Practical Optimization*, Academic Press, London, UK, 1981.

See Also

\ | lsqnonneg | quadprog | optimtool

How To

• "Linear Least Squares"

Purpose Solve nonlinear least-squares (nonlinear data-fitting) problems

Equation Solves nonlinear least-squares curve fitting problems of the form

$$\min_x \|f(x)\|_2^2 = \min_x (f_1(x)^2 + f_2(x)^2 + \dots + f_n(x)^2)$$

with optional lower and upper bounds lb and ub on the components of x . x , lb , and ub can be vectors or matrices; see “Matrix Arguments” on page 2-32.

Syntax

```
x = lsqnonlin(fun,x0)
x = lsqnonlin(fun,x0,lb,ub)
x = lsqnonlin(fun,x0,lb,ub,options)
x = lsqnonlin(problem)
[x,resnorm] = lsqnonlin(...)
[x,resnorm,residual] = lsqnonlin(...)
[x,resnorm,residual,exitflag] = lsqnonlin(...)
[x,resnorm,residual,exitflag,output] = lsqnonlin(...)
[x,resnorm,residual,exitflag,output,lambda] = lsqnonlin(...)
[x,resnorm,residual,exitflag,output,lambda,
jacobian] = lsqnonlin(...)
```

Description

`lsqnonlin` solves nonlinear least-squares problems, including nonlinear data-fitting problems.

Rather than compute the value $\|f(x)\|_2^2$ (the sum of squares), `lsqnonlin` requires the user-defined function to compute the *vector*-valued function

$$f(x) = \begin{bmatrix} f_1(x) \\ f_2(x) \\ \vdots \\ f_n(x) \end{bmatrix}$$

Then, in vector terms, you can restate this optimization problem as

$$\min_x \|f(x)\|_2^2 = \min_x (f_1(x)^2 + f_2(x)^2 + \dots + f_n(x)^2)$$

where x is a vector or matrix and $f(x)$ is a function that returns a vector or matrix value. For details of matrix values, see “Matrix Arguments” on page 2-32.

$x = \text{lsqnonlin}(\text{fun}, x0)$ starts at the point $x0$ and finds a minimum of the sum of squares of the functions described in `fun`. `fun` should return a vector of values and not the sum of squares of the values. (The algorithm implicitly computes the sum of squares of the components of `fun(x)`.)

Note “Passing Extra Parameters” on page 2-53 explains how to pass extra parameters to the vector function f , if necessary.

$x = \text{lsqnonlin}(\text{fun}, x0, lb, ub)$ defines a set of lower and upper bounds on the design variables in x , so that the solution is always in the range $lb \leq x \leq ub$.

$x = \text{lsqnonlin}(\text{fun}, x0, lb, ub, \text{options})$ minimizes with the optimization options specified in `options`. Use `optimoptions` to set these options. Pass empty matrices for `lb` and `ub` if no bounds exist.

$x = \text{lsqnonlin}(\text{problem})$ finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-235.

Create the `problem` structure by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x, resnorm] = lsqnonlin(...)` returns the value of the squared 2-norm of the residual at x : `sum(fun(x).^2)`.

`[x, resnorm, residual] = lsqnonlin(...)` returns the value of the residual `fun(x)` at the solution x .

`[x, resnorm, residual, exitflag] = lsqnonlin(...)` returns a value `exitflag` that describes the exit condition.

`[x,resnorm,residual,exitflag,output] = lsqnonlin(...)` returns a structure output that contains information about the optimization.

`[x,resnorm,residual,exitflag,output,lambda] = lsqnonlin(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

`[x,resnorm,residual,exitflag,output,lambda, jacobian] = lsqnonlin(...)` returns the Jacobian of `fun` at the solution `x`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the outputs `resnorm` and `residual` are `[]`.

Components of `x0` that violate the bounds `lb` `x` `ub` are reset to the interior of the box defined by the bounds. Components that respect the bounds are not changed.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into `lsqnonlin`. This section provides function-specific details for `fun`, `options`, and `problem`:

`fun` The function whose sum of squares is minimized. `fun` is a function that accepts a vector `x` and returns a vector `F`, the objective functions evaluated at `x`. The function `fun` can be specified as a function handle to a file:

```
x = lsqnonlin(@myfun,x0)
```

where `myfun` is a MATLAB function such as

```
function F = myfun(x)
F = ...                    % Compute function values at x
```

`fun` can also be a function handle for an anonymous function.

```
x = lsqnonlin(@(x)sin(x.*x),x0);
```

If the user-defined values for x and F are matrices, they are converted to a vector using linear indexing.

Note The sum of squares should not be formed explicitly. Instead, your function should return a vector of function values. See “Examples” on page 10-245.

If the Jacobian can also be computed *and* the Jacobian option is 'on', set by

```
options = optimoptions('lsqnonlin','Jacobian','on')
```

the function `fun` must return, in a second output argument, the Jacobian value J , a matrix, at x . By checking the value of `nargout`, the function can avoid computing J when `fun` is called with only one output argument (in the case where the optimization algorithm only needs the value of F but not J).

```
function [F,J] = myfun(x)
F = ...           % Objective function values at x
if nargout > 1   % Two output arguments
    J = ...       % Jacobian of the function evaluated at x
end
```

If `fun` returns a vector (matrix) of m components and x has length n , where n is the length of `x0`, the Jacobian J is an m -by- n matrix where $J(i,j)$ is the partial derivative of $F(i)$ with respect to $x(j)$. (The Jacobian J is the transpose of the gradient of F .)

`options` “Options” on page 10-238 provides the function-specific details for the `options` values.

problem	objective	Objective function
	x0	Initial point for x
	lb	Vector of lower bounds
	ub	Vector of upper bounds
	solver	'lsqnonlin'
	options	Options created with optimoptions

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by `lsqnonlin`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated:
1	Function converged to a solution <code>x</code> .
2	Change in <code>x</code> was less than the specified tolerance.
3	Change in the residual was less than the specified tolerance.
4	Magnitude of search direction was smaller than the specified tolerance.
0	Number of iterations exceeded <code>options.MaxIter</code> or number of function evaluations exceeded <code>options.MaxFunEvals</code> .
-1	Output function terminated the algorithm.

	-2	Problem is infeasible: the bounds lb and ub are inconsistent.
	-4	Line search could not sufficiently decrease the residual along the current search direction.
lambda		Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields are
	lower	Lower bounds lb
	upper	Upper bounds ub
output		Structure containing information about the optimization. The fields of the structure are
	firstorderopt	Measure of first-order optimality (trust-region-reflective algorithm, [] for others)
	iterations	Number of iterations taken
	funcCount	The number of function evaluations
	cgiterations	Total number of PCG iterations (trust-region-reflective algorithm, [] for others)
	stepsize	Final displacement in x (Levenberg-Marquardt algorithm)
	algorithm	Optimization algorithm used
	message	Exit message

Options

Optimization options. Set or change options using the `optimoptions` function. Some options apply to all algorithms, some are only relevant when you are using the trust-region-reflective algorithm, and others are only relevant when you are using the Levenberg-Marquardt

algorithm. See “Optimization Options Reference” on page 9-7 for detailed information.

Algorithm Options

Both algorithms use the following options:

Algorithm	<p>Choose between 'trust-region-reflective' (default) and 'levenberg-marquardt'. Set the initial Levenberg-Marquardt parameter λ by setting Algorithm to a cell array such as {'levenberg-marquardt', .005}. The default $\lambda = 0.01$.</p> <p>The Algorithm option specifies a preference for which algorithm to use. It is only a preference, because certain conditions must be met to use each algorithm. For the trust-region-reflective algorithm, the nonlinear system of equations cannot be underdetermined; that is, the number of equations (the number of elements of F returned by fun) must be at least as many as the length of x. The Levenberg-Marquardt algorithm does not handle bound constraints. For more information on choosing the algorithm, see “Choosing the Algorithm” on page 2-7.</p>
DerivativeCheck	<p>Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives. The choices are 'on' or the default 'off'.</p>
Diagnostics	<p>Display diagnostic information about the function to be minimized or solved. The choices are 'on' or the default 'off'.</p>
DiffMaxChange	<p>Maximum change in variables for finite-difference gradients (a positive scalar). The default is Inf.</p>

DiffMinChange Minimum change in variables for finite-difference gradients (a positive scalar). The default is 0.

Display Level of display:

- 'off' or 'none' displays no output.
- 'iter' displays output at each iteration, and gives the default exit message.
- 'iter-detailed' displays output at each iteration, and gives the technical exit message.
- 'final' (default) displays just the final output, and gives the default exit message.
- 'final-detailed' displays just the final output, and gives the technical exit message.

FinDiffRelStep Scalar or vector step size factor. When you set **FinDiffRelStep** to a vector **v**, forward finite differences **delta** are

```
delta =  
v.*sign(x).*max(abs(x),TypicalX);
```

and central finite differences are

```
delta = v.*max(abs(x),TypicalX);
```

Scalar **FinDiffRelStep** expands to a vector. The default is $\sqrt{\text{eps}}$ for forward finite differences, and $\text{eps}^{1/3}$ for central finite differences.

FinDiffType Finite differences, used to estimate gradients, are either 'forward' (default), or 'central' (centered). 'central' takes twice as many function evaluations, but should be more accurate.

The algorithm is careful to obey bounds when estimating both types of finite differences. So, for example, it could take a backward, rather than a forward, difference to avoid evaluating at a point outside bounds.

FunValCheck	Check whether function values are valid. 'on' displays an error when the function returns a value that is complex, Inf, or NaN. The default 'off' displays no error.
Jacobian	If 'on', lsqnonlin uses a user-defined Jacobian (defined in fun), or Jacobian information (when using JacobMult), for the objective function. If 'off' (default), lsqnonlin approximates the Jacobian using finite differences.
MaxFunEvals	Maximum number of function evaluations allowed, a positive integer. The default is 100*numberOfVariables.
MaxIter	Maximum number of iterations allowed, a positive integer. The default is 400.
OutputFcn	Specify one or more user-defined functions that an optimization function calls at each iteration, either as a function handle or as a cell array of function handles. The default is none ([]). See “Output Function” on page 9-21.
PlotFcns	Plots various measures of progress while the algorithm executes, select from predefined plots or write your own. Pass a function handle or a cell array of function handles. The default is none ([]): <ul style="list-style-type: none">• @optimplotx plots the current point.• @optimplotfunccount plots the function count.• @optimplotfval plots the function value.• @optimplotresnorm plots the norm of the residuals.• @optimplotstepsize plots the step size.• @optimplotfirstorderopt plots the first-order optimality measure.

	For information on writing a custom plot function, see “Plot Functions” on page 9-30.
TolFun	Termination tolerance on the function value, a positive scalar. The default is $1e-6$.
TolX	Termination tolerance on x , a positive scalar. The default is $1e-6$.
TypicalX	Typical x values. The number of elements in TypicalX is equal to the number of elements in x_0 , the starting point. The default value is <code>ones(numberofvariables,1)</code> . <code>lsqnonlin</code> uses TypicalX for scaling finite differences for gradient estimation.

Trust-Region-Reflective Algorithm Only

The trust-region-reflective algorithm uses the following options:

JacobMult	Function handle for Jacobian multiply function. For large-scale structured problems, this function computes the Jacobian matrix product $J*Y$, $J'*Y$, or $J'*(J*Y)$ without actually forming J . The function is of the form
-----------	---

$$W = \text{jmfun}(Jinfo, Y, flag)$$

where $Jinfo$ contains the matrix used to compute $J*Y$ (or $J'*Y$, or $J'*(J*Y)$). The first argument $Jinfo$ must be the same as the second argument returned by the objective function `fun`, for example, by

$$[F, Jinfo] = \text{fun}(x)$$

Y is a matrix that has the same number of rows as there are dimensions in the problem. `flag` determines which product to compute:

- If `flag == 0` then $W = J'*(J*Y)$.
- If `flag > 0` then $W = J*Y$.
- If `flag < 0` then $W = J'*Y$.

In each case, J is not formed explicitly. `lsqnonlin` uses `Jinfo` to compute the preconditioner. See “Passing Extra Parameters” on page 2-53 for information on how to supply values for any additional parameters `jmfun` needs.

Note 'Jacobian' must be set to 'on' for `lsqnonlin` to pass `Jinfo` from `fun` to `jmfun`.

See “Minimization with Dense Structured Hessian, Linear Equalities” on page 6-80 and “Jacobian Multiply Function with Linear Least Squares” on page 6-224 for similar examples.

JacobPattern

Sparsity pattern of the Jacobian for finite differencing. Set `JacobPattern(i,j) = 1` when `fun(i)` depends on `x(j)`. Otherwise, set `JacobPattern(i,j) = 0`. In other words, `JacobPattern(i,j) = 1` when you can have $\partial \text{fun}(i) / \partial x(j) \neq 0$.

Use `JacobPattern` when it is inconvenient to compute the Jacobian matrix J in `fun`, though you can determine (say, by inspection) when `fun(i)` depends on `x(j)`. `lsqnonlin` can approximate J via sparse finite differences when you give `JacobPattern`.

In the worst case, if the structure is unknown, do not set `JacobPattern`. The default behavior

is as if `JacobPattern` is a dense matrix of ones. Then `lsqnonlin` computes a full finite-difference approximation in each iteration. This can be very expensive for large problems, so it is usually better to determine the sparsity structure.

- | | |
|-------------------------------|---|
| <code>MaxPCGIter</code> | Maximum number of PCG (preconditioned conjugate gradient) iterations, a positive scalar. The default is $\max(1, \text{numberOfVariables}/2)$. For more information, see “Algorithms” on page 10-217. |
| <code>PrecondBandWidth</code> | Upper bandwidth of preconditioner for PCG, a nonnegative integer. The default <code>PrecondBandWidth</code> is <code>Inf</code> , which means a direct factorization (Cholesky) is used rather than the conjugate gradients (CG). The direct factorization is computationally more expensive than CG, but produces a better quality step towards the solution. Set <code>PrecondBandWidth</code> to 0 for diagonal preconditioning (upper bandwidth of 0). For some problems, an intermediate bandwidth reduces the number of PCG iterations. |
| <code>TolPCG</code> | Termination tolerance on the PCG iteration, a positive scalar. The default is 0.1. |

Levenberg-Marquardt Algorithm Only

The Levenberg-Marquardt algorithm uses the following options:

- | | |
|---------------------------|---|
| <code>ScaleProblem</code> | 'Jacobian' can sometimes improve the convergence of a poorly scaled problem; the default is 'none'. |
|---------------------------|---|

Examples

Find x that minimizes

$$\sum_{k=1}^{10} (2 + 2k - e^{kx_1} - e^{kx_2})^2,$$

starting at the point $x = [0.3, 0.4]$.

Because `lsqnonlin` assumes that the sum of squares is *not* explicitly formed in the user-defined function, the function passed to `lsqnonlin` should instead compute the vector-valued function

$$F_k(x) = 2 + 2k - e^{kx_1} - e^{kx_2},$$

for $k = 1$ to 10 (that is, F should have 10 components).

First, write a file to compute the 10-component vector F .

```
function F = myfun(x)
k = 1:10;
F = 2 + 2*k - exp(k*x(1)) - exp(k*x(2));
```

Next, invoke an optimization routine.

```
x0 = [0.3 0.4] % Starting guess
[x,resnorm] = lsqnonlin(@myfun,x0); % Invoke optimizer
```

After about 24 function evaluations, this example gives the solution

```
x,resnorm
x =
    0.2578    0.2578

resnorm =
    124.3622
```

Algorithms**Trust-Region-Reflective Optimization**

By default, `lsqnonlin` chooses the trust-region-reflective algorithm. This algorithm is a subspace trust-region method and is based on

the interior-reflective Newton method described in [1] and [2]. Each iteration involves the approximate solution of a large linear system using the method of preconditioned conjugate gradients (PCG). See “Trust-Region-Reflective Least Squares” on page 6-201, and in particular “Large Scale Nonlinear Least Squares” on page 6-203.

Levenberg-Marquardt Optimization

If you set the Algorithm option to 'levenberg-marquardt' using optimoptions, lsqnonlin uses the Levenberg-Marquardt method [4], [5], and [6]. See “Levenberg-Marquardt Method” on page 6-205.

Diagnostics

Trust-Region-Reflective Optimization

The trust-region-reflective method does not allow equal upper and lower bounds. For example, if $lb(2) == ub(2)$, lsqin gives the error

Equal upper and lower bounds not permitted.

(lsqnonlin does not handle equality constraints, which is another way to formulate equal bounds. If equality constraints are present, use fmincon, fminimax, or fgoalattain for alternative formulations where equality constraints can be included.)

Limitations

The function to be minimized must be continuous. lsqnonlin might only give local solutions.

lsqnonlin can solve complex-valued problems directly with the levenberg-marquardt algorithm. However, this algorithm does not accept bound constraints. For a complex problem with bound constraints, split the variables into real and imaginary parts, and use the trust-region-reflective algorithm. See “Fit a Model to Complex-Valued Data” on page 6-231.

Trust-Region-Reflective Optimization

The trust-region-reflective algorithm for lsqnonlin does not solve underdetermined systems; it requires that the number of equations, i.e., the row dimension of F , be at least as great as the number of variables.

In the underdetermined case, the Levenberg-Marquardt algorithm is used instead.

The preconditioner computation used in the preconditioned conjugate gradient part of the trust-region-reflective method forms $J^T J$ (where J is the Jacobian matrix) before computing the preconditioner; therefore, a row of J with many nonzeros, which results in a nearly dense product $J^T J$, can lead to a costly solution process for large problems.

If components of x have no upper (or lower) bounds, `lsqnonlin` prefers that the corresponding components of `ub` (or `lb`) be set to `inf` (or `-inf` for lower bounds) as opposed to an arbitrary but very large positive (or negative for lower bounds) number.

Trust-Region-Reflective Problem Coverage and Requirements

For Large Problems

- Provide sparsity structure of the Jacobian or compute the Jacobian in `fun`.
- The Jacobian should be sparse.

Levenberg-Marquardt Optimization

The Levenberg-Marquardt algorithm does not handle bound constraints.

Since the trust-region-reflective algorithm does not handle underdetermined systems and the Levenberg-Marquardt does not handle bound constraints, problems with both these characteristics cannot be solved by `lsqnonlin`.

References

- [1] Coleman, T.F. and Y. Li, “An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds,” *SIAM Journal on Optimization*, Vol. 6, pp. 418–445, 1996.
- [2] Coleman, T.F. and Y. Li, “On the Convergence of Reflective Newton Methods for Large-Scale Nonlinear Minimization Subject to Bounds,” *Mathematical Programming*, Vol. 67, Number 2, pp. 189-224, 1994.

[3] Dennis, J.E., Jr., “Nonlinear Least-Squares,” *State of the Art in Numerical Analysis*, ed. D. Jacobs, Academic Press, pp. 269–312, 1977.

[4] Levenberg, K., “A Method for the Solution of Certain Problems in Least-Squares,” *Quarterly Applied Math.* 2, pp. 164–168, 1944.

[5] Marquardt, D., “An Algorithm for Least-Squares Estimation of Nonlinear Parameters,” *SIAM Journal Applied Math.*, Vol. 11, pp. 431–441, 1963.

[6] Moré, J.J., “The Levenberg-Marquardt Algorithm: Implementation and Theory,” *Numerical Analysis*, ed. G. A. Watson, *Lecture Notes in Mathematics* 630, Springer Verlag, pp. 105–116, 1977.

See Also

`lsqcurvefit` | `lsqlin` | `optimoptions` | `optimtool`

How To

- `function_handle`
- “Nonlinear Least Squares (Curve Fitting)”

Purpose Solve nonnegative least-squares constraint problem

Equation Solves nonnegative least-squares curve fitting problems of the form

$$\min_x \|C \cdot x - d\|_2^2, \text{ where } x \geq 0.$$

Syntax

```
x = lsqnonneg(C,d)
x = lsqnonneg(C,d,options)
x = lsqnonneg(problem)
[x,resnorm] = lsqnonneg(...)
[x,resnorm,residual] = lsqnonneg(...)
[x,resnorm,residual,exitflag] = lsqnonneg(...)
[x,resnorm,residual,exitflag,output] = lsqnonneg(...)
[x,resnorm,residual,exitflag,output,lambda] = lsqnonneg(...)
```

Description

`x = lsqnonneg(C,d)` returns the vector `x` that minimizes `norm(C*x-d)` subject to `x ≥ 0`. `C` and `d` must be real.

`x = lsqnonneg(C,d,options)` minimizes with the optimization options specified in the structure `options`. Use `optimset` to set these options.

`x = lsqnonneg(problem)` finds the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-250.

Create the structure `problem` by exporting a problem from Optimization app, as described in “Exporting Your Work” on page 5-13.

`[x,resnorm] = lsqnonneg(...)` returns the value of the squared 2-norm of the residual, `norm(C*x-d)^2`.

`[x,resnorm,residual] = lsqnonneg(...)` returns the residual `d-C*x`.

`[x,resnorm,residual,exitflag] = lsqnonneg(...)` returns a value `exitflag` that describes the exit condition of `lsqnonneg`.

`[x,resnorm,residual,exitflag,output] = lsqnonneg(...)` returns a structure `output` that contains information about the optimization.

[x,resnorm,residual,exitflag,output,lambda] =
lsqnonneg(...) returns the Lagrange multipliers in the vector lambda.

Input Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments passed into lsqnonneg. This section provides function-specific details for options and problem:

options		Use optimset to set or change the values of these fields in the options structure, options. See “Optimization Options Reference” on page 9-7 for detailed information.
	Display	Level of display: <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'final' displays just the final output.• 'notify' (default) displays output only if the function does not converge.
	TolX	Termination tolerance on x, a positive scalar. The default is $10 \cdot \text{eps} \cdot \text{norm}(C,1) \cdot \text{length}(C)$.
problem	C	Matrix
	d	Vector
	x0	Initial point for x
	solver	'lsqnonneg'
	options	Options structure created with optimset

Output Arguments

“Function Arguments” on page 9-2 contains general descriptions of arguments returned by lsqnonneg. This section provides function-specific details for exitflag, lambda, and output:

<code>exitflag</code>	Integer identifying the reason the algorithm terminated. The following lists the values of <code>exitflag</code> and the corresponding reasons the algorithm terminated.						
	<table> <tr> <td>1</td> <td>Function converged to a solution x.</td> </tr> <tr> <td>0</td> <td>Number of iterations exceeded <code>options.MaxIter</code>.</td> </tr> </table>	1	Function converged to a solution x .	0	Number of iterations exceeded <code>options.MaxIter</code> .		
1	Function converged to a solution x .						
0	Number of iterations exceeded <code>options.MaxIter</code> .						
<code>lambda</code>	Vector containing the Lagrange multipliers: <code>lambda(i)</code> 0 when $x(i)$ is (approximately) 0, and <code>lambda(i)</code> is (approximately) 0 when $x(i) > 0$.						
<code>output</code>	Structure containing information about the optimization. The fields are <table> <tr> <td><code>iterations</code></td> <td>Number of iterations taken</td> </tr> <tr> <td><code>algorithm</code></td> <td>'active-set'</td> </tr> <tr> <td><code>message</code></td> <td>Exit message</td> </tr> </table>	<code>iterations</code>	Number of iterations taken	<code>algorithm</code>	'active-set'	<code>message</code>	Exit message
<code>iterations</code>	Number of iterations taken						
<code>algorithm</code>	'active-set'						
<code>message</code>	Exit message						

Examples

Compare the unconstrained least-squares solution to the `lsqnonneg` solution for a 4-by-2 problem.

```
C = [
    0.0372    0.2869
    0.6861    0.7071
    0.6233    0.6245
    0.6344    0.6170];
d = [
    0.8587
    0.1781
    0.0747
    0.8405];

[C\d, lsqnonneg(C,d)]
```

```
ans =  
    -2.5627    0  
     3.1108    0.6929  
  
[norm(C*(C\d)-d), norm(C*lsqnonneg(C,d)-d)]  
ans =  
    0.6674    0.9118
```

The solution from `lsqnonneg` does not fit as well as the least-squares solution. However, the nonnegative least-squares solution has no negative components.

Algorithms

`lsqnonneg` uses the algorithm described in [1]. The algorithm starts with a set of possible basis vectors and computes the associated dual vector `lambda`. It then selects the basis vector corresponding to the maximum value in `lambda` in order to swap it out of the basis in exchange for another possible candidate. This continues until `lambda` ≤ 0 .

Notes

The nonnegative least-squares problem is a subset of the constrained linear least-squares problem. Thus, when `C` has more rows than columns (i.e., the system is overdetermined),

```
[x,resnorm,residual,exitflag,output,lambda] = ...  
    lsqnonneg(C,d)
```

is equivalent to

```
[m,n] = size(C);  
[x,resnorm,residual,exitflag,output,lambda_lsqli] = ...  
    lsqli(C,d,-eye(n,n),zeros(n,1));
```

except that `lambda` = `-lambda_lsqli.ineqli`.

For problems greater than order 20, `lsqli` might be faster than `lsqnonneg`; otherwise `lsqnonneg` is generally more efficient.

References

[1] Lawson, C.L. and R.J. Hanson, *Solving Least-Squares Problems*, Prentice-Hall, Chapter 23, p. 161, 1974.

See Also

\ | lsqlin | optimset | optimtool

optimget

Purpose Optimization options values

Syntax

```
val = optimget(options,'param')  
val = optimget(options,'param',default)
```

Description `val = optimget(options,'param')` returns the value of the specified option in the optimization options structure `options`. You need to type only enough leading characters to define the option name uniquely. Case is ignored for option names.

`val = optimget(options,'param',default)` returns `default` if the specified option is not defined in the optimization options structure `options`. Note that this form of the function is used primarily by other optimization functions.

Examples This statement returns the value of the `Display` option in the structure called `my_options`.

```
val = optimget(my_options,'Display')
```

This statement returns the value of the `Display` option in the structure called `my_options` (as in the previous example) except that if the `Display` option is not defined, it returns the value `'final'`.

```
optnew = optimget(my_options,'Display','final');
```

See Also `optimset`

Purpose

Create optimization options

Syntax

```
options = optimoptions(SolverName)
options = optimoptions(SolverName,Name,Value)

options = optimoptions(olddoptions,Name,Value)

options = optimoptions(SolverName,olddoptions)
```

Description

`options = optimoptions(SolverName)` returns a set of default options for the `SolverName` solver.

`options = optimoptions(SolverName,Name,Value)` returns options with the named parameters altered with the specified values.

`options = optimoptions(olddoptions,Name,Value)` returns a copy of `olddoptions` with the named parameters altered with the specified values.

`options = optimoptions(SolverName,olddoptions)` returns default options for the `SolverName` solver, and copies the applicable options in `olddoptions` to `options`.

Input Arguments

SolverName - Solver name

string or function handle

Solver name, specified as a string or function handle.

Example: 'fmincon'

Example: @fmincon

Data Types

char | function_handle

olddoptions - Options

options created using `optimoptions`

optimoptions

Options, specified as an options object. The `optimoptions` function creates options objects.

Example: `oldoptions = optimoptions(@fminunc)`

Name-Value Pair Arguments

Specify optional comma-separated pairs of `Name`, `Value` arguments. `Name` is the argument name and `Value` is the corresponding value. `Name` must appear inside single quotes (' '). You can specify several name and value pair arguments in any order as `Name1,Value1,...,NameN,ValueN`.

Example:

`optimoptions(@fmincon, 'Display', 'iter', 'TolFun', 1e-10)` sets `fmincon` options to have iterative display, and to have a `TolFun` tolerance of `1e-10`.

For relevant name-value pairs, look at the options table for your solver:

- `bintprog` options
- `fgoalattain` options
- `fmincon` options
- `fminimax` options
- `fminunc` options
- `fseminf` options
- `fsolve` options
- `intlinprog` options
- `ktrlink` options
- `linprog` options
- `lsqcurvefit` options
- `lsqlin` options
- `lsqnonlin` options

- quadprog options

Output Arguments

options - Options object

options object

Options object, containing options for the SolverName solver.

Examples

Create Default Options

Create default options for the fmincon solver.

```
options = optimoptions('fmincon')
```

optimoptions

options =

[fmincon](#) options:

Options used by current Algorithm ('interior-point'):

([Other available algorithms](#): 'active-set', 'sqp', 'trust-region-reflective')

Set by user:

No options set by user.

Default:

```
Algorithm: 'interior-point'
AlwaysHonorConstraints: 'bounds'
DerivativeCheck: 'off'
Diagnostics: 'off'
DiffMaxChange: Inf
DiffMinChange: 0
Display: 'final'
FinDiffRelStep: 'sqrt(eps)'
FinDiffType: 'forward'
FunValCheck: 'off'
GradConstr: 'off'
GradObj: 'off'
HessFcn: []
Hessian: 'bfgs'
HessMult: []
InitBarrierParam: 0.1000
InitTrustRegionRadius: 'sqrt(numberOfVariables)'
MaxFunEvals: 3000
MaxIter: 1000
MaxProjCGIter: '2*(numberOfVariables-numberOfEqualities)'
ObjectiveLimit: -1.0000e+20
OutputFcn: []
PlotFcns: []
ScaleProblem: 'none'
SubproblemAlgorithm: 'ldl-factorization'
TolCon: 1.0000e-06
TolFun: 1.0000e-06
TolProjCG: 0.0100
TolProjCGAbs: 1.0000e-10
TolX: 1.0000e-10
```


Create Nondefault Options

Create nondefault options for the `fmincon` solver.

```
options = optimoptions(@fmincon,'Algorithm','sqp','MaxIter',1500)
```

optimoptions

options =

[fmincon](#) options:

Options used by current Algorithm ('sqp'):

([Other available algorithms](#): 'active-set', 'interior-point', 'trust-region-refle

Set by user:

Algorithm: 'sqp'

MaxIter: 1500

Default:

DerivativeCheck: 'off'

Diagnostics: 'off'

DiffMaxChange: Inf

DiffMinChange: 0

Display: 'final'

FinDiffRelStep: 1.4901e-08

FinDiffType: 'forward'

FunValCheck: 'off'

GradConstr: 'off'

GradObj: 'off'

MaxFunEvals: '100*numberOfVariables'

ObjectiveLimit: -1.0000e+20

OutputFcn: []

PlotFcns: []

ScaleProblem: 'none'

TolCon: 1.0000e-06

TolFun: 1.0000e-06

TolX: 1.0000e-06

TypicalX: 'ones(numberOfVariables,1)'

UseParallel: 'never'

[Show options](#) not used by current Algorithm ('sqp')

Update Options

Update existing options with new values.

Create options for the lsqnonlin solver.

```
oldoptions = optimoptions(@lsqnonlin, 'Algorithm', 'levenberg-marquardt
```

optimoptions

```
oldoptions =
```

```
lsqnonlin options:
```

```
Options used by current Algorithm ('levenberg-marquardt'):  
(Other available algorithms: 'trust-region-reflective')
```

```
Set by user:
```

```
Algorithm: 'levenberg-marquardt'  
MaxFunEvals: 1500
```

```
Default:
```

```
DerivativeCheck: 'off'  
Diagnostics: 'off'  
DiffMaxChange: Inf  
DiffMinChange: 0  
Display: 'final'  
FinDiffRelStep: 1.4901e-08  
FinDiffType: 'forward'  
FunValCheck: 'off'  
Jacobian: 'off'  
MaxIter: 400  
OutputFcn: []  
PlotFcns: []  
ScaleProblem: 'none'  
TolFun: 1.0000e-06  
TolX: 1.0000e-06  
TypicalX: 'ones(numberOfVariables,1)'
```

```
Show options not used by current Algorithm ('levenberg-marquardt')
```

Increase MaxFunEvals to 2000.

```
options = optimoptions(oldoptions,'MaxFunEvals',2000)

options =

lsgmonlin options:

Options used by current Algorithm ('levenberg-marquardt'):
(Other available algorithms: 'trust-region-reflective')

Set by user:
  Algorithm: 'levenberg-marquardt'
  MaxFunEvals: 2000

Default:
  DerivativeCheck: 'off'
    Diagnostics: 'off'
  DiffMaxChange: Inf
  DiffMinChange: 0
    Display: 'final'
  FinDiffRelStep: 1.4901e-08
  FinDiffType: 'forward'
  FunValCheck: 'off'
    Jacobian: 'off'
    MaxIter: 400
  OutputFcn: []
  PlotFcns: []
  ScaleProblem: 'none'
    TolFun: 1.0000e-06
    TolX: 1.0000e-06
  TypicalX: 'ones(numberOfVariables,1)'

Show options not used by current Algorithm ('levenberg-marquardt')
```

optimoptions

Use Dot Notation to Update Options

Update existing options with new values by using dot notation.

Create options for the lsqnonlin solver.

```
options = optimoptions(@lsqnonlin,'Algorithm','levenberg-marquardt','MaxF
```

options =

[lsgmonlin](#) options:

Options used by current Algorithm ('levenberg-marquardt'):
([Other available algorithms](#): 'trust-region-reflective')

Set by user:

Algorithm: 'levenberg-marquardt'
MaxFunEvals: 1500

Default:

DerivativeCheck: 'off'
Diagnostics: 'off'
DiffMaxChange: Inf
DiffMinChange: 0
Display: 'final'
FinDiffRelStep: 1.4901e-08
FinDiffType: 'forward'
FunValCheck: 'off'
Jacobian: 'off'
MaxIter: 400
OutputFcn: []
PlotFcns: []
ScaleProblem: 'none'
TolFun: 1.0000e-06
TolX: 1.0000e-06
TypicalX: 'ones(numberOfVariables,1)'

[Show options](#) not used by current Algorithm ('levenberg-marquardt')

Increase MaxFunEvals to 2000 by using dot notation.

```
options.MaxFunEvals = 2000
```

optimoptions

options =

[lsqmonlin](#) options:

Options used by current Algorithm ('levenberg-marquardt'):
([Other available algorithms](#): 'trust-region-reflective')

Set by user:

Algorithm: 'levenberg-marquardt'
MaxFunEvals: 2000

Default:

DerivativeCheck: 'off'
Diagnostics: 'off'
DiffMaxChange: Inf
DiffMinChange: 0
Display: 'final'
FinDiffRelStep: 1.4901e-08
FinDiffType: 'forward'
FunValCheck: 'off'
Jacobian: 'off'
MaxIter: 400
OutputFcn: []
PlotFcns: []
ScaleProblem: 'none'
TolFun: 1.0000e-06
TolX: 1.0000e-06
TypicalX: 'ones(numberOfVariables,1)'

[Show options](#) not used by current Algorithm ('levenberg-marquardt')

Copy Options to Another Solver

Transfer nondefault options for the `fmincon` solver to options for the `fminunc` solver.

Create nondefault options for the `fmincon` solver.

```
oldoptions = optimoptions(@fmincon,'Algorithm','sqp','MaxIter',1500)
```

optimoptions

```
oldoptions =
```

```
fmincon options:
```

```
Options used by current Algorithm ('sqp'):
```

```
(Other available algorithms: 'active-set', 'interior-point', 'trust-region-refle
```

```
Set by user:
```

```
Algorithm: 'sqp'
```

```
MaxIter: 1500
```

```
Default:
```

```
DerivativeCheck: 'off'
```

```
  Diagnostics: 'off'
```

```
DiffMaxChange: Inf
```

```
DiffMinChange: 0
```

```
  Display: 'final'
```

```
FinDiffRelStep: 1.4901e-08
```

```
  FinDiffType: 'forward'
```

```
FunValCheck: 'off'
```

```
  GradConstr: 'off'
```

```
  GradObj: 'off'
```

```
MaxFunEvals: '100*numberOfVariables'
```

```
ObjectiveLimit: -1.0000e+20
```

```
  OutputFcn: []
```

```
  PlotFcns: []
```

```
ScaleProblem: 'none'
```

```
  TolCon: 1.0000e-06
```

```
  TolFun: 1.0000e-06
```

```
  TolX: 1.0000e-06
```

```
  TypicalX: 'ones(numberOfVariables,1)'
```

```
UseParallel: 'never'
```

```
Show options not used by current Algorithm ('sqp')
```

Transfer the applicable options to the `fminunc` solver.

```
options = optimoptions(@fminunc,oldoptions)
```

optimoptions

```
options =
```

```
fminunc options:
```

```
Options used by current Algorithm ('trust-region'):  
(Other available algorithms: 'quasi-newton')
```

```
Set by user:
```

```
MaxIter: 1500
```

```
Default:
```

```
    Algorithm: 'trust-region'  
DerivativeCheck: 'off'  
    Diagnostics: 'off'  
    DiffMaxChange: Inf  
    DiffMinChange: 0  
    Display: 'final'  
FinDiffRelStep: 1.4901e-08  
    FinDiffType: 'forward'  
    FunValCheck: 'off'  
    GradObj: 'off'  
    HessMult: []  
    HessPattern: 'sparse(ones(numberofvariables))'  
    Hessian: 'off'  
    MaxFunEvals: '100*numberofvariables'  
    MaxPCGIter: 'max(1,floor(numberofvariables/2))'  
    OutputFcn: []  
    PlotFcns: []  
PrecondBandWidth: 0  
    TolFun: 1.0000e-06  
    TolPCG: 0.1000  
    TolX: 1.0000e-06  
    TypicalX: 'ones(numberofvariables,1)'
```

```
Show options not used by current Algorithm ('trust-region')
```

Alternative Functionality

App

You can set and modify options using the Optimization app (optimtool).

See Also

optimset | optimtool

Concepts

- “Set Options”
- “Optimization App” on page 5-2

optimset

Purpose Create or edit optimization options structure

Syntax

```
options = optimset('param1',value1,'param2',value2,...)
optimset
options = optimset
options = optimset(optimfun)
options = optimset(oldopts,'param1',value1,...)
options = optimset(oldopts,newopts)
```

Description `options = optimset('param1',value1,'param2',value2,...)` creates an optimization options structure called `options`, in which the specified options (`param`) have specified values. Any unspecified options are set to `[]` (options with value `[]` indicate to use the default value for that option when you pass `options` to the optimization function). It is sufficient to type only enough leading characters to define the option name uniquely. Case is ignored for option names.

`optimset` with no input or output arguments displays a complete list of options with their valid values.

`options = optimset` (with no input arguments) creates an options structure `options` where all fields are set to `[]`.

`options = optimset(optimfun)` creates an options structure `options` with all option names and default values relevant to the optimization function `optimfun`.

`options = optimset(oldopts,'param1',value1,...)` creates a copy of `oldopts`, modifying the specified options with the specified values.

`options = optimset(oldopts,newopts)` combines an existing options structure, `oldopts`, with a new options structure, `newopts`. Any options in `newopts` with nonempty values overwrite the corresponding old options in `oldopts`.

Options For more information about individual options, including their default values, see the reference pages for the optimization functions. “Optimization Options Reference” on page 9-7 provides descriptions of optimization options and which functions use them.

Use the command `optimset(@solver)` or the equivalent `optimset solver` to see the default values of relevant optimization options for a solver. Some solvers do not have a default value, since the default depends on the algorithm. For example, the default value of the `maxIter` option in the `fmincon` solver is 400 for the trust-region-reflective algorithm, but is 1000 for the interior-point algorithm.

You can also see the default values of all relevant options in the Optimization app. To see the options:

- 1 Start the Optimization app, e.g., with the `optimtool` command.
- 2 Choose the solver from the **Solver** menu.
- 3 Choose the algorithm, if applicable, from the **Algorithm** menu.
- 4 Read off the default values within the **Options** pane.

Examples

This statement creates an optimization options structure called `options` in which the `Display` option is set to `'iter'` and the `TolFun` option is set to `1e-8`.

```
options = optimset('Display','iter','TolFun',1e-8)
```

This statement makes a copy of the options structure called `options`, changing the value of the `TolX` option and storing new values in `optnew`.

```
optnew = optimset(options,'TolX',1e-4);
```

This statement returns an optimization options structure `options` that contains all the option names and default values relevant to the function `fminbnd`.

```
options = optimset('fminbnd')
```

If you only want to see the default values for `fminbnd`, you can simply type

```
optimset fminbnd
```

optimset

or equivalently

```
optimset('fminbnd')
```

See Also

[optimget](#) | [optimoptions](#) | [optimtool](#)

Purpose Select solver and optimization options, run problems

Syntax

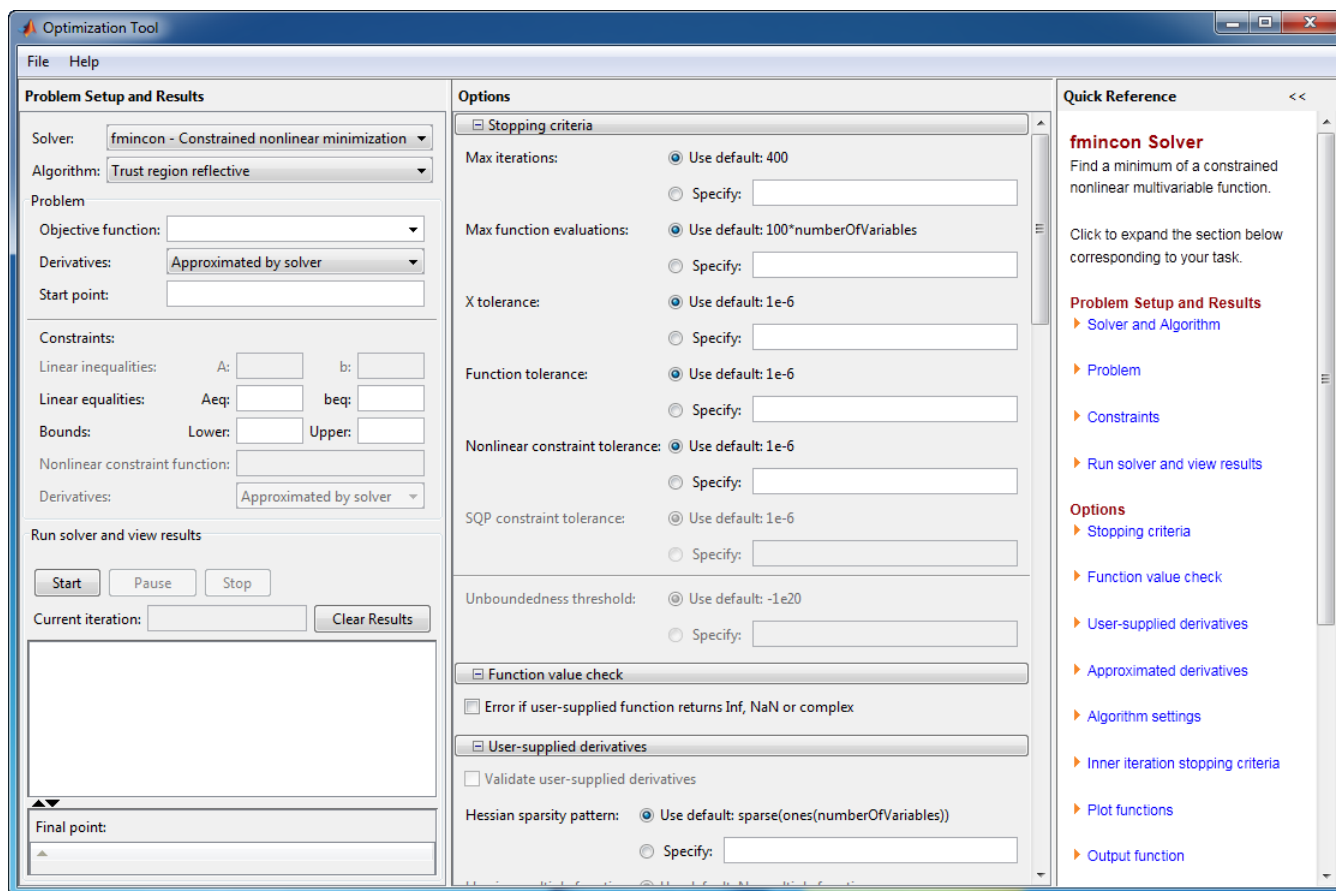
```
optimtool  
optimtool(optstruct)  
optimtool('solver')
```

Description `optimtool` opens the Optimization app. Use the Optimization app to select a solver, optimization options, and run problems. See “Optimization App” on page 5-2 for a complete description of the Optimization app.

The Optimization app can be used to run any Optimization Toolbox solver, and any Global Optimization Toolbox solver except `GlobalSearch` and `MultiStart`. Results can be exported to a file or to the MATLAB workspace as a structure.

`optimtool(optstruct)` starts the Optimization app and loads `optstruct`. `optstruct` can be either optimization options or an optimization problem structure. Create optimization options with the `optimoptions` or `optimset` function, or by using the export option from the Optimization app. Create a problem structure by exporting the problem from the Optimization app to the MATLAB workspace. If you have Global Optimization Toolbox, you can create a problem structure with the `createOptimProblem` function.

`optimtool('solver')` starts the Optimization app with the specified solver, identified as a string, and the corresponding default options and problem fields. All Optimization Toolbox and Global Optimization Toolbox solvers are valid inputs to the `optimtool` function, except `GlobalSearch` and `MultiStart`.



See Also

[optimoptions](#) | [optimset](#)

How To

- “Optimization App” on page 5-2
- “Solve a Constrained Nonlinear Problem” on page 1-3

Purpose

Quadratic programming

Syntax

```

x = quadprog(H,f)
x = quadprog(H,f,A,b)
x = quadprog(H,f,A,b,Aeq,beq)
x = quadprog(H,f,A,b,Aeq,beq,lb,ub)
x = quadprog(H,f,A,b,Aeq,beq,lb,ub,x0)
x = quadprog(H,f,A,b,Aeq,beq,lb,ub,x0,options)
x = quadprog(problem)
[x,fval] = quadprog(H,f,...)
[x,fval,exitflag] = quadprog(H,f,...)
[x,fval,exitflag,output] = quadprog(H,f,...)
[x,fval,exitflag,output,lambda] = quadprog(H,f,...)

```

Description

Finds a minimum for a problem specified by

$$\min_x \frac{1}{2} x^T H x + f^T x \text{ such that } \begin{cases} A \cdot x \leq b, \\ Aeq \cdot x = beq, \\ lb \leq x \leq ub. \end{cases}$$

H , A , and Aeq are matrices, and f , b , beq , lb , ub , and x are vectors.

f , lb , and ub can be passed as vectors or matrices; see “Matrix Arguments” on page 2-32.

$x = \text{quadprog}(H,f)$ returns a vector x that minimizes $1/2 * x' * H * x + f' * x$. H must be positive definite for the problem to have a finite minimum.

$x = \text{quadprog}(H,f,A,b)$ minimizes $1/2 * x' * H * x + f' * x$ subject to the restrictions $A * x \leq b$. A is a matrix of doubles, and b is a vector of doubles.

$x = \text{quadprog}(H,f,A,b,Aeq,beq)$ solves the preceding problem subject to the additional restrictions $Aeq * x = beq$. Aeq is a matrix of doubles, and beq is a vector of doubles. If no inequalities exist, set $A = []$ and $b = []$.

`x = quadprog(H,f,A,b,Aeq,beq,lb,ub)` solves the preceding problem subject to the additional restrictions $lb \leq x \leq ub$. `lb` and `ub` are vectors of doubles, and the restrictions hold for each `x` component. If no equalities exist, set `Aeq = []` and `beq = []`.

Note If the specified input bounds for a problem are inconsistent, the output `x` is `x0` and the output `fval` is `[]`.

`quadprog` resets components of `x0` that violate the bounds $lb \leq x \leq ub$ to the interior of the box defined by the bounds. `quadprog` does not change components that respect the bounds.

`x = quadprog(H,f,A,b,Aeq,beq,lb,ub,x0)` solves the preceding problem starting from the vector `x0`. If no bounds exist, set `lb = []` and `ub = []`. Some `quadprog` algorithms ignore `x0`, see “Input Arguments” on page 10-279.

`x = quadprog(H,f,A,b,Aeq,beq,lb,ub,x0,options)` solves the preceding problem using the optimization options specified in `options`. Use `optimoptions` to create `options`. If you do not want to give an initial point, set `x0 = []`.

`x = quadprog(problem)` returns the minimum for `problem`, where `problem` is a structure described in “Input Arguments” on page 10-279. Create `problem` by exporting a problem using the Optimization app; see “Exporting Your Work” on page 5-13.

`[x,fval] = quadprog(H,f,...)` returns the value of the objective function at `x`:

$$fval = 0.5*x'*H*x + f'*x$$

`[x,fval,exitflag] = quadprog(H,f,...)` `exitflag`, a scalar that describes the exit condition of `quadprog`.

`[x,fval,exitflag,output] = quadprog(H,f,...)` `output`, a structure that contains information about the optimization.

`[x,fval,exitflag,output,lambda] = quadprog(H,f,...)` lambda, a structure whose fields contain the Lagrange multipliers at the solution x.

Input Arguments

H

Symmetric matrix of doubles. Represents the quadratic in the expression $1/2*x'*H*x + f'*x$.

f

Vector of doubles. Represents the linear term in the expression $1/2*x'*H*x + f'*x$.

A

Matrix of doubles. Represents the linear coefficients in the constraints $A*x \leq b$.

b

Vector of doubles. Represents the constant vector in the constraints $A*x \leq b$.

Aeq

Matrix of doubles. Represents the linear coefficients in the constraints $Aeq*x = beq$.

beq

Vector of doubles. Represents the constant vector in the constraints $Aeq*x = beq$.

lb

Vector of doubles. Represents the lower bounds elementwise in $lb \leq x \leq ub$.

ub

Vector of doubles. Represents the upper bounds elementwise in $lb \leq x \leq ub$.

x0

Vector of doubles. Optional. The initial point for some quadprog algorithms:

- active-set
- trust-region-reflective when there are only bound constraints

If you do not give x0, quadprog sets all components of x0 to a point in the interior of the box defined by the bounds. quadprog ignores x0 for the interior-point-convex algorithm, and for the trust-region-reflective algorithm with equality constraints.

options

Options created using optimoptions or the Optimization app.

All Algorithms

Algorithm

Choose the algorithm:

- 'interior-point-convex' (default)
- 'trust-region-reflective'
- 'active-set'

The 'trust-region-reflective' algorithm handles problems with only bounds, or only linear equality constraints, but not both. The 'interior-point-convex' algorithm handles only convex problems. For details, see “Choosing the Algorithm” on page 2-7.

Diagnostics

Display diagnostic information about the function to be minimized or solved. The choices are 'on' or 'off' (default).

All Algorithms (Continued)

Display	<p>Level of display returned to the command line.</p> <ul style="list-style-type: none">• 'off' or 'none' displays no output.• 'final' displays just the final output (default). <p>The 'interior-point-convex' algorithm allows additional values:</p> <ul style="list-style-type: none">• 'iter' gives iterative display.• 'iter-detailed' gives iterative display with a detailed exit message.• 'final-detailed' displays just the final output, with a detailed exit message.
MaxIter	<p>Maximum number of iterations allowed, a positive integer.</p> <ul style="list-style-type: none">• For a 'trust-region-reflective' equality-constrained problem, the default value is $2 * (\text{numberOfVariables} - \text{numberOfEqualities})$.• For all other algorithms and problems, the default value is 200.

All Algorithms Except active-set

TolFun	Termination tolerance on the function value, a positive scalar. <ul style="list-style-type: none">• For a 'trust-region-reflective' equality-constrained problem, the default value is 1e-6.• For a 'trust-region-reflective' bound-constrained problem, the default value is 100*eps, about 2.2204e-14.• For 'interior-point-convex', the default value is 1e-8.
TolX	Termination tolerance on x, a positive scalar. <ul style="list-style-type: none">• For 'trust-region-reflective', the default value is 100*eps, about 2.2204e-14.• For 'interior-point-convex', the default value is 1e-8.

trust-region-reflective Algorithm Only

HessMult	Function handle for a Hessian multiply function. For large-scale structured problems, this function computes the Hessian matrix product $H*Y$ without actually forming H . The function has the form $W = \text{hmfun}(\text{Hinfo}, Y)$ where Hinfo and possibly some additional parameters contain the matrices used to compute $H*Y$.
----------	--

trust-region-reflective Algorithm Only (Continued)

	See “Quadratic Minimization with Dense, Structured Hessian” on page 6-138 for an example that uses this option.
MaxPCGIter	Maximum number of PCG (preconditioned conjugate gradient) iterations, a positive scalar. The default is $\max(1, \text{floor}(\text{numberOfVariables}/2))$. For more information, see “Preconditioned Conjugate Gradient Method” on page 6-128.
PrecondBandWidth	Upper bandwidth of the preconditioner for PCG, a nonnegative integer. By default, quadprog uses diagonal preconditioning (upper bandwidth 0). For some problems, increasing the bandwidth reduces the number of PCG iterations. Setting PrecondBandWidth to Inf uses a direct factorization (Cholesky) rather than the conjugate gradients (CG). The direct factorization is computationally more expensive than CG, but produces a better quality step toward the solution.
TolPCG	Termination tolerance on the PCG iteration, a positive scalar. The default is 0.1.
TypicalX	Typical x values. The number of elements in TypicalX equals the number of elements in x0, the starting point. The default value is $\text{ones}(\text{numberOfVariables}, 1)$. quadprog uses TypicalX internally for scaling. TypicalX has an effect only when x has unbounded components, and when a TypicalX value for an unbounded component exceeds 1.

interior-point-convex Algorithm Only

TolCon Tolerance on the constraint violation, a positive scalar. The default is $1e-8$.

problem

Structure encapsulating the quadprog inputs and options:

H	Symmetric matrix in $1/2*x'*H*x$
f	Vector in linear term $f'*x$
Aineq	Matrix in linear inequality constraints $Aineq*x \leq bineq$
bineq	Vector in linear inequality constraints $Aineq*x \leq bineq$
Aeq	Matrix in linear equality constraints $Aeq*x = beq$
beq	Vector in linear equality constraints $Aeq*x = beq$
lb	Vector of lower bounds
ub	Vector of upper bounds
x0	Initial point for x
solver	'quadprog'
options	Options created using optimoptions or the Optimization app

Output Arguments

x
Vector that minimizes $1/2*x'*H*x + f'*x$ subject to all bounds and linear constraints. x can be a local minimum for nonconvex problems. For convex problems, x is a global minimum. For more information, see “Local vs. Global Optima” on page 4-27.

fval

Value of $1/2*x'*H*x + f'*x$ at the solution x , a double.

exitflag

Integer identifying the reason the algorithm terminated. The following lists the values of `exitflag` and the corresponding reasons the algorithm terminated:

All Algorithms

1	Function converged to the solution x .
0	Number of iterations exceeded <code>options.MaxIter</code> .
-2	Problem is infeasible.
-3	Problem is unbounded.

interior-point-convex Algorithm

-6	Nonconvex problem detected.
----	-----------------------------

trust-region-reflective Algorithm

3	Change in the objective function value was smaller than <code>options.TolFun</code> .
-4	Current search direction was not a direction of descent. No further progress could be made.

active-set Algorithm

4	Local minimizer was found.
-7	Magnitude of search direction became too small. No further progress could be made. The problem is ill-posed or badly conditioned.

output

Structure containing information about the optimization. The fields are:

iterations	Number of iterations taken
algorithm	Optimization algorithm used
cgiterations	Total number of PCG iterations (trust-region-reflective algorithm only)
constrviolation	Maximum of constraint functions
firstorderopt	Measure of first-order optimality
message	Exit message

lambda

Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields are:

lower	Lower bounds lb
upper	Upper bounds ub
ineqlin	Linear inequalities
eqlin	Linear equalities

For details, see “Lagrange Multiplier Structures” on page 3-27.

Examples

Solve a simple quadratic programming problem: find values of x that minimize

$$f(x) = \frac{1}{2}x_1^2 + x_2^2 - x_1x_2 - 2x_1 - 6x_2,$$

subject to

$$\begin{aligned}x_1 + x_2 &\leq 2 \\ -x_1 + 2x_2 &\leq 2 \\ 2x_1 + x_2 &\leq 3 \\ 0 &\leq x_1, 0 \leq x_2.\end{aligned}$$

In matrix notation this is

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

where

$$H = \begin{bmatrix} 1 & -1 \\ -1 & 2 \end{bmatrix}, \quad f = \begin{bmatrix} -2 \\ -6 \end{bmatrix}, \quad x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

1 Enter the coefficient matrices:

```
H = [1 -1; -1 2];  
f = [-2; -6];  
A = [1 1; -1 2; 2 1];  
b = [2; 2; 3];  
lb = zeros(2,1);
```

2 Set the options to use the 'active-set' algorithm with no display:

```
opts =  
optimoptions('quadprog','Algorithm','active-set','Display','off');
```

3 Call quadprog:

```
[x,fval,exitflag,output,lambda] = ...  
quadprog(H,f,A,b,[],[],lb,[],[],opts);
```

4 Examine the final point, function value, and exit flag:

```
x,fval,exitflag  
  
x =  
    0.6667  
    1.3333  
  
fval =  
   -8.2222
```

```
exitflag =  
    1
```

- 5** An exit flag of 1 means the result is a local minimum. Because H is a positive definite matrix, this problem is convex, so the minimum is a global minimum. You can see H is positive definite by noting all its eigenvalues are positive:

```
eig(H)  
ans =  
    0.3820  
    2.6180
```

Use the 'interior-point-convex' algorithm to solve a sparse quadratic program.

- 1** Generate a sparse symmetric matrix for the quadratic form:

```
v = sparse([1, -.25, 0, 0, 0, 0, 0, -.25]);  
H = gallery('circul', v);
```

- 2** Include the linear term for the problem:

```
f = -4:3;
```

- 3** Include the constraint that the sum of the terms in the solution x must be less than -2:

```
A = ones(1, 8); b = -2;
```

- 4** Set options to use the 'interior-point-convex' algorithm and iterative display:

```
opts = optimoptions('quadprog', ...  
    'Algorithm', 'interior-point-convex', 'Display', 'iter');
```

5 Run the quadprog solver and observe the iterations:

```
[x fval eflag output lambda] =
quadprog(H,f,A,b,[],[],[],[],[],opts);
```

Iter	f(x)	Feasibility	First-order optimality	Total relative error
0	-2.000000e+000	1.000e+001	4.500e+000	1.200e+001
1	-2.630486e+001	0.000e+000	9.465e-002	9.465e-002
2	-2.639877e+001	0.000e+000	3.914e-005	3.914e-005
3	-2.639881e+001	0.000e+000	3.069e-015	6.883e-015

Minimum found that satisfies the constraints.

Optimization completed because the objective function is non-decreasing in feasible directions, to within the default value of the function tolerance, and constraints are satisfied to within the default value of the constraint tolerance.

6 Examine the solution:

```
fval,eflag
```

```
fval =
    -26.3988
```

```
eflag =
     1
```

For the 'interior-point-convex' algorithm, an exit flag of 1 means the result is a global minimum.

Algorithms

trust-region-reflective

The 'trust-region-reflective' algorithm is a subspace trust-region method based on the interior-reflective Newton method described in [1]. Each iteration involves the approximate solution of a large linear system using the method of preconditioned conjugate gradients (PCG).

For more information, see “trust-region-reflective quadprog Algorithm” on page 6-126.

active-set

quadprog uses an active set method, which is also a projection method, similar to that described in [2]. It finds an initial feasible solution by first solving a linear programming problem. For more information, see “active-set quadprog Algorithm” on page 6-131.

interior-point-convex

The 'interior-point-convex' algorithm attempts to follow a path that is strictly inside the constraints. It uses a presolve module to remove redundancies, and to simplify the problem by solving for components that are straightforward. For more information, see “interior-point-convex quadprog Algorithm” on page 6-121.

References

- [1] Coleman, T.F. and Y. Li, “A Reflective Newton Method for Minimizing a Quadratic Function Subject to Bounds on Some of the Variables,” *SIAM Journal on Optimization*, Vol. 6, Number 4, pp. 1040–1058, 1996.
- [2] Gill, P. E., W. Murray, and M. H. Wright, *Practical Optimization*, Academic Press, London, UK, 1981.
- [3] Gould, N. and P. L. Toint. “Preprocessing for quadratic programming.” *Math. Programming*, Series B, Vol. 100, pp. 95–132, 2004.

Alternatives

You can use the Optimization app for quadratic programming. Enter `optimtool` at the MATLAB command line, and choose the `quadprog` - Quadratic programming solver. For more information, see “Optimization App” on page 5-2.

See Also

`linprog` | `lsqlin` | `optimoptions`

How To

- “Optimization Problem Setup”

- “Optimization Results”
- “Optimization App” on page 5-2
- “Quadratic Programming”