

Optimization Toolbox

For Use with MATLAB®

Computation

Visualization

Programming

User's Guide

Version 2

The
MATH
WORKS

How to Contact The MathWorks:



www.mathworks.com
comp.soft-sys.matlab

Web
Newsgroup



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

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



508-647-7000

Phone



508-647-7001

Fax



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

Mail

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

Optimization Toolbox User's Guide

© COPYRIGHT 1990 - 2001 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 or for the federal government of the United States. By accepting delivery of the Program, the government hereby agrees that this software qualifies as "commercial" computer software within the meaning of FAR Part 12.212, DFARS Part 227.7202-1, DFARS Part 227.7202-3, DFARS Part 252.227-7013, and DFARS Part 252.227-7014. The terms and conditions of The MathWorks, Inc. Software License Agreement shall pertain to the government's use and disclosure of the Program and Documentation, and shall supersede any conflicting contractual terms or conditions. If this license fails to meet the government's minimum needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to MathWorks.

MATLAB, Simulink, Stateflow, Handle Graphics, and Real-Time Workshop are registered trademarks, and Target Language Compiler is a trademark of The MathWorks, Inc.

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

Printing 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 Revised for Version 2.1.1 (Release 12.1) (online only)

Preface

What Is the Optimization Toolbox?	viii
New Features in Version 2.1	ix
Toolbox Speed	ix
Function Handles	ix
Large Structured Problems	ix
Large Problems that Are Not Well-Scaled	x
Using This Guide	xi
Acknowledgments	xii
Related Products	xiii
Configuration Information	xv
Technical Conventions	xvi
Matrix, Vector, and Scalar Notation	xvi
Typographical Conventions	xvii

Tutorial

1

Overview	1-2
Introduction	1-4
Problems Covered by the Toolbox	1-4
Using the Optimization Functions	1-6

Examples that Use Standard Algorithms	1-8
Unconstrained Example	1-9
Nonlinear Inequality Constrained Example	1-10
Constrained Example with Bounds	1-12
Constrained Example with Gradients	1-13
Gradient Check: Analytic Versus Numeric	1-15
Equality Constrained Example	1-16
Maximization	1-17
Greater-Than-Zero Constraints	1-17
Additional Arguments: Avoiding Global Variables	1-18
Multiobjective Examples	1-18
Large-Scale Examples	1-30
Problems Covered by Large-Scale Methods	1-31
Nonlinear Equations with Jacobian	1-34
Nonlinear Equations with Jacobian Sparsity Pattern	1-36
Nonlinear Least Squares with Full Jacobian Sparsity Pattern	1-38
Nonlinear Minimization with Gradient and Hessian	1-40
Nonlinear Minimization with Gradient and Hessian Sparsity Pattern	1-41
Nonlinear Minimization with Bound Constraints and Banded Preconditioner	1-43
Nonlinear Minimization with Equality Constraints	1-47
Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints	1-48
Quadratic Minimization with Bound Constraints	1-52
Quadratic Minimization with a Dense but Structured Hessian	1-54
Linear Least Squares with Bound Constraints	1-58
Linear Programming with Equalities and Inequalities	1-59
Linear Programming with Dense Columns in the Equalities	1-60
Default Parameter Settings	1-63
Changing the Default Settings	1-63
Displaying Iterative Output	1-65
Output Headings: Medium-Scale Algorithms	1-65
Output Headings: Large-Scale Algorithms	1-67
Optimization of Inline Objects Instead of M-Files	1-70

Typical Problems and How to Deal with Them	1-72
Converting Your Code to Version 2 Syntax	1-76
Using optimset and optimget	1-77
New Calling Sequences	1-77
Example of Converting from constr to fmincon	1-85
Selected Bibliography	1-88

Standard Algorithms

2

Overview	2-2
Optimization Overview	2-4
Unconstrained Optimization	2-5
Quasi-Newton Methods	2-7
Line Search	2-9
Quasi-Newton Implementation	2-12
Hessian Update	2-12
Line Search Procedures	2-12
Least Squares Optimization	2-19
Gauss-Newton Method	2-20
Levenberg-Marquardt Method	2-22
Nonlinear Least Squares Implementation	2-24
Gauss-Newton Implementation	2-24
Levenberg-Marquardt Implementation	2-24
Constrained Optimization	2-26
Sequential Quadratic Programming (SQP)	2-27
Quadratic Programming (QP) Subproblem	2-28

SQP Implementation	2-30
Updating the Hessian Matrix	2-30
Quadratic Programming Solution	2-31
Line Search and Merit Function	2-35
Multiobjective Optimization	2-36
Introduction	2-36
Goal Attainment Method	2-42
Algorithm Improvements for Goal Attainment Method	2-43
Selected Bibliography	2-46

Large-Scale Algorithms

3

Overview	3-2
Trust Region Methods for Nonlinear Minimization	3-3
Preconditioned Conjugate Gradients	3-6
Linearly Constrained Problems	3-8
Linear Equality Constraints	3-8
Box Constraints	3-8
Nonlinear Least Squares	3-11
Quadratic Programming	3-12
Linear Least Squares	3-13
Large-Scale Linear Programming	3-14
Main Algorithm	3-14
Preprocessing	3-17
Selected Bibliography	3-18

Functions by Category	4-3
Minimization	4-3
Equation Solving	4-3
Least Squares (Curve Fitting)	4-4
Utility	4-4
Demos of Large-Scale Methods	4-4
Demos of Medium-Scale Methods	4-5
Function Arguments	4-6
Optimization Options Parameters	4-11
fgoalattain	4-16
fminbnd	4-28
fmincon	4-32
fminimax	4-45
fminsearch	4-54
fminunc	4-58
fseminf	4-68
fsolve	4-81
fzero	4-92
linprog	4-96
lsqcurvefit	4-104
lsqlin	4-115
lsqnonlin	4-123
lsqnonneg	4-134
optimget	4-138
optimset	4-139
quadprog	4-143

Preface

What Is the Optimization Toolbox?	viii
New Features in Version 2.1	ix
Toolbox Speed	ix
Function Handles	ix
Large Structured Problems	ix
Large Problems that Are Not Well-Scaled	x
Using This Guide	xi
Acknowledgmentsxii
Related Products	xiii
Configuration Informationxv
Technical Conventions	xvi
Matrix, Vector, and Scalar Notation	xvi
Typographical Conventions	xvii

What Is the Optimization Toolbox?

The Optimization Toolbox is a collection of functions that extend the capability of the MATLAB® numeric computing environment. The toolbox includes routines for many types of optimization including:

- Unconstrained nonlinear minimization
- Constrained nonlinear minimization, including goal attainment problems, minimax problems, and semi-infinite minimization problems
- Quadratic and linear programming
- Nonlinear least squares and curve-fitting
- Nonlinear system of equation solving
- Constrained linear least squares
- Sparse and structured large-scale problems

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

```
type function_name
```

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

This Preface also provides information about:

- New features in Version 2.1 and 2.1.1
- Using this guide, which includes a summary of the chapters
- Acknowledgements
- Related MathWorks products
- Checking if the Optimization Toolbox is installed on your system
- Technical and typographical conventions used in this manual

New Features in Version 2.1

Version 2.1 of the Optimization Toolbox offers a number of advances over previous versions:

- Overall toolbox speed is improved.
- Functions that take a function as an argument now support the use of function handles.
- Large structured problems are supported.
- Support is improved in large-scale methods for problems that are not well-scaled.

Toolbox Speed

By improving the speed of `optimset` and `optimget`, the overall speed of the toolbox is improved. Applications that call these functions repeatedly should exhibit improved execution time.

Function Handles

You can now use the new `@` syntax to specify a function handle for any argument that expects a function. As in earlier versions of the toolbox, you can also specify the function as a string or as an inline object.

For more information about function handles, see:

- The `function_handle (@)`, `func2str`, and `str2func` reference pages
- The “Function Handles” section of “Programming and Data Types” in the MATLAB documentation

Large Structured Problems

The functions `fmincon`, `fminunc`, `fsolve`, `lsqcurvefit`, `lsqlin`, `lsqnonlin`, and `quadprog` now support solving large structured problems, i.e., problems that have large dense Hessian or Jacobian matrices that you do not want to form explicitly, but for which Hessian-matrix (or Jacobian-matrix) products are efficient to compute.

Two new options parameters, `HessMult` and `JacobMult`, provide access to this new feature. For more information, see “Optimization Options Parameters” on page 4-11 and the individual reference pages.

Large Problems that Are Not Well-Scaled

In Version 2.1.1, large-scale finite differencing is improved numerically to handle cases when an optimization problem is not well-scaled. These changes potentially improve the speed and accuracy of results when using the large-scale versions of `lsqnonlin`, `lsqcurvefit`, `fsolve`, `fmincon` and `fminunc`, particularly if the objective function is not well-scaled.

Using This Guide

This guide has the following chapters:

- The “Tutorial” chapter shows you how to solve a variety of different optimization problems. It includes a section that highlights large-scale problems. This chapter also provides information on how to use the toolbox functions in conjunction with Simulink using multiobjective optimization. Other sections include information about changing default parameters and using inline objects.
- The “Standard Algorithms” and “Large-Scale Algorithms” chapters describe the algorithms used by the optimization functions. “Standard Algorithms” describes the problem formulations and algorithms for the medium-scale algorithms. “Large-Scale Algorithms” focuses on algorithms used to solve large sparse or structured problems.
- The “Function Reference” chapter provides a detailed reference description of each toolbox function. Reference descriptions include the function’s syntax, a description of the different calling sequences available, and detailed information about arguments to the function, including relevant optimization options parameters. Reference descriptions may also include examples, a summary of the function’s algorithms, and references to additional reading material.

Acknowledgments

The MathWorks would like to acknowledge these contributors:

Thomas F. Coleman researched and contributed the large scale 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 Computer Science and Applied Mathematics at Cornell University. He is Director of the Cornell Theory Center and the Cornell Computational Finance Institute. Dr. Coleman is Chair of the SIAM Activity Group on Optimization, and a member of the Editorial Boards of *Applied Mathematics Letters*, *SIAM Journal of Scientific Computing*, *Computational Optimization and Applications*, *Communications on Applied Nonlinear Analysis*, and *Mathematical Modeling and Scientific Computing*.

Dr. Coleman has published 4 books and over 70 technical papers in the areas of continuous optimization and computational methods and tools for large-scale problems.

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

Dr. Zhang is Associate Professor of Computational and Applied Mathematics and is on the faculty of the Keck Center for Computational Biology at Rice University. He is on the Editorial Board of *SIAM Journal on Optimization*, and Associate Editor of *Journal of Optimization: Theory and Applications*.

Dr. Zhang has published over 40 technical papers in the areas of interior-point methods for linear programming and computation mathematical programming.

Related Products

The MathWorks provides several products that are relevant to the kinds of tasks you can perform with the Optimization Toolbox. For more information about any of these products, see either:

- The online documentation for that product, if it is installed or if you are reading the documentation from the CD
- The MathWorks Web site, at <http://www.mathworks.com>; see the “products” section

Note The toolboxes listed below all include functions that extend MATLAB’s capabilities. The blocksets all include blocks that extend Simulink’s capabilities.

Product	Description
Data Acquisition Toolbox	MATLAB functions for direct access to live, measured data from MATLAB
Database Toolbox	Tool for connecting to, and interacting with, most ODBC/JDBC databases from within MATLAB
Financial Time Series Toolbox	Tool for analyzing time series data in the financial markets
Financial Toolbox	MATLAB functions for quantitative financial modeling and analytic prototyping
Frequency Domain System Identification Toolbox	Tool for frequency domain model identification and validation
GARCH Toolbox	MATLAB functions for univariate Generalized Autoregressive Conditional Heteroskedasticity (GARCH) volatility modeling

Product	Description
LMI Control Toolbox	Convex optimization algorithms for solving linear matrix inequalities (LMI), with application to robust control, multi-objective control, and gain scheduling
Neural Network Toolbox	Comprehensive environment for neural network research, design, and simulation within MATLAB
Nonlinear Control Design Blockset	Simulink block libraries that provide a time-domain-based optimization approach to system design; automatically tunes parameters based on user-defined time-domain performance constraints
Signal Processing Toolbox	Tool for algorithm development, signal and linear system analysis, and time-series data modeling
Simulink	Interactive, graphical environment for modeling, simulating, and prototyping dynamic systems
Spline Toolbox	Tool for the construction and use of piecewise polynomial functions
Statistics Toolbox	Tool for analyzing historical data, modeling systems, developing statistical algorithms, and learning and teaching statistics
Symbolic Math Toolbox	Tool for seamless integration of symbolic mathematics and variable precision arithmetic based on Maple V
System Identification Toolbox	Tool for building accurate, simplified models of complex systems from noisy time-series data

Configuration Information

To determine if the Optimization Toolbox is installed on your system, type this command at the MATLAB prompt.

```
ver
```

When you enter this command, MATLAB displays information about the version of MATLAB you are running, including a list of all toolboxes installed on your system and their version numbers.

If the Optimization Toolbox is not installed, check the *Installation Guide* for your platform for instructions on how to install it.

Note For the most up-to-date information about system requirements, see the system requirements page, available in the products area at the MathWorks Web site (<http://www.mathworks.com>).

Technical Conventions

Matrix, Vector, and Scalar Notation

Upper-case letters such as A are used to denote matrices. Lower-case letters such as x are used to denote vectors, except where noted that it is a scalar.

For functions, the notation differs slightly to follow the usual conventions in optimization. For vector functions, we use an upper-case letter such as F in $F(x)$. A function that returns a scalar value is denoted with a lower-case letter such as f in $f(x)$.

Typographical Conventions

This manual uses some or all of these conventions.

Item	Convention Used	Example
Example code	Monospace font	To assign the value 5 to A, enter <code>A = 5</code>
Function names/syntax	Monospace font	The <code>cos</code> function finds the cosine of each array element. Syntax line example is <code>MLGetVar ML_var_name</code>
Keys	Boldface with an initial capital letter	Press the R eturn key.
Literal strings (in syntax descriptions in reference chapters)	Monospace bold for literals	<code>f = freqspace(n, 'whole')</code>
Mathematical expressions	<i>Italics</i> for variables Standard text font for functions, operators, and constants	This vector represents the polynomial $p = x^2 + 2x + 3$
MATLAB output	Monospace font	MATLAB responds with <code>A =</code> <code>5</code>
Menu titles, menu items, dialog boxes, and controls	Boldface with an initial capital letter	Choose the F ile menu.
New terms	<i>Italics</i>	An <i>array</i> is an ordered collection of information.
Omitted input arguments	(...) ellipsis denotes all of the input/output arguments from preceding syntaxes.	<code>[c,ia,ib] = union(...)</code>
String variables (from a finite list)	<i>Monospace italics</i>	<code>sysc = d2c(sysd, 'method')</code>

Tutorial

Overview	1-2
Introduction	1-4
Examples that Use Standard Algorithms	1-8
Large-Scale Examples	1-30
Default Parameter Settings	1-63
Displaying Iterative Output	1-65
Optimization of Inline Objects Instead of M-Files	1-70
Typical Problems and How to Deal with Them	1-72
Converting Your Code to Version 2 Syntax	1-76
Selected Bibliography	1-88

Overview

This “Tutorial” provides information on how to use the toolbox functions. It also provides examples for solving different optimization problems. It is organized into these major topics.

Introduction

Summarizes, in tabular form, the functions available for minimization, equation solving, and solving least squares or data fitting problems. It also provides basic guidelines for using the optimization routines and introduces the algorithms and line-search strategies that are available for solving medium- and large-scale problems.

Examples that Use Standard Algorithms

Presents *medium-scale* algorithms through a selection of minimization examples. These examples include unconstrained and constrained problems, as well as problems with and without user-supplied gradients. This section also discusses maximization, greater-than-zero constraints, passing additional arguments, and multiobjective examples.

Large-Scale Examples

Presents *large-scale* algorithms through a selection of large-scale examples. These examples include specifying sparsity structures, and preconditioners, as well as unconstrained and constrained problems.

Default Parameter Settings

Describes the use of default parameter settings and tells you how to change them. It also tells you how to determine which parameters are used by a specified function, and provides examples of setting some commonly used parameters.

Displaying Iterative Output

Describes the column headings used in the iterative output of both medium-scale and large-scale algorithms.

Optimization of Inline Objects Instead of M-Files

Tells you how to represent a mathematical function at the command line by creating an `inline` object from a string expression.

Typical Problems and How to Deal with Them

Provides tips to help you improve solutions found using the optimization functions, improve efficiency of the algorithms, overcome common difficulties, and transform problems that are not typically in standard form.

Converting Your Code to Version 2 Syntax

For each function, compares a Version 1.5 call to the equivalent Version 2 call. This section also describes the Version 2 calling sequences and provides a detailed example of converting from `constr` to its replacement `fmincon`.

Selected Bibliography

Lists published materials that support concepts implemented in the Statistics Toolbox.

Introduction

Optimization concerns the minimization or maximization of functions. The Optimization Toolbox consists of functions that perform minimization (or maximization) on general nonlinear functions. Functions for nonlinear equation solving and least squares (data-fitting) problems are also provided.

This introduction includes the following sections:

- Problems Covered by the Toolbox
- Using the Optimization Functions

Problems Covered by the Toolbox

The tables below show the functions available for minimization, equation solving, and solving least squares or data fitting problems.

Note The following tables list the types of problems in order of increasing complexity.

Table 1-1: Minimization

Type	Notation	Function
Scalar Minimization	$\min_a f(a)$ such that $a_1 < a < a_2$	fminbnd
Unconstrained Minimization	$\min_x f(x)$	fminunc, fminsearch
Linear Programming	$\min_x f^T x$ such that $A \cdot x \leq b, \text{ Aeq} \cdot x = \text{beq}, l \leq x \leq u$	linprog
Quadratic Programming	$\min_x \frac{1}{2} x^T H x + f^T x$ such that $A \cdot x \leq b, \text{ Aeq} \cdot x = \text{beq}, l \leq x \leq u$	quadprog

Table 1-1: Minimization (Continued)

Type	Notation	Function
Constrained Minimization	$\min_x f(x) \text{ such that}$ $c(x) \leq 0, \text{ ceq}(x) = 0$ $A \cdot x \leq b, \text{ Aeq} \cdot x = \text{beq}, \quad l \leq x \leq u$	fmincon
Goal Attainment	$\min_{x, \gamma} \gamma \text{ such that}$ $F(x) - w\gamma \leq \text{goal}$ $c(x) \leq 0, \text{ ceq}(x) = 0$ $A \cdot x \leq b, \text{ Aeq} \cdot x = \text{beq}, \quad l \leq x \leq u$	fgoalattain
Minimax	$\min_x \max_{\{F_i\}} \{F_i(x)\} \text{ such that}$ $c(x) \leq 0, \text{ ceq}(x) = 0$ $A \cdot x \leq b, \text{ Aeq} \cdot x = \text{beq}, \quad l \leq x \leq u$	fminimax
Semi-Infinite Minimization	$\min_x f(x) \text{ such that}$ $K(x, w) \leq 0 \text{ for all } w$ $c(x) \leq 0, \text{ ceq}(x) = 0$ $A \cdot x \leq b, \text{ Aeq} \cdot x = \text{beq}, \quad l \leq x \leq u$	fseminf

Table 1-2: Equation Solving

Type	Notation	Function
Linear Equations	$C \cdot x = d, n \text{ equations}, n \text{ variables}$	\ (slash)
Nonlinear Equation of One Variable	$f(a) = 0$	fzero

Table 1-2: Equation Solving (Continued)

Type	Notation	Function
Nonlinear Equations	$F(x) = 0$, n equations, n variables	<code>fsolve</code>

Table 1-3: Least Squares (Curve Fitting)

Type	Notation	Function
Linear Least Squares	$\min_x \ C \cdot x - d\ _2^2$, m equations, n variables	<code>\</code> (slash)
Nonnegative Linear Least Squares	$\min_x \ C \cdot x - d\ _2^2$ such that $x \geq 0$	<code>lsqnonneg</code>
Constrained Linear Least Squares	$\min_x \ C \cdot x - d\ _2^2$ such that $A \cdot x \leq b$, $Aeq \cdot x = beq$, $l \leq x \leq u$	<code>lsqlin</code>
Nonlinear Least Squares	$\min_x \frac{1}{2} \ F(x)\ _2^2 = \frac{1}{2} \sum_i F_i(x)^2$ such that $l \leq x \leq u$	<code>lsqnonlin</code>
Nonlinear Curve Fitting	$\min_x \frac{1}{2} \ F(x, xdata) - ydata\ _2^2$ such that $l \leq x \leq u$	<code>lsqcurvefit</code>

Using the Optimization Functions

Most of these optimization routines require the definition of an M-file containing the function to be minimized, i.e., the objective function. Alternatively, an inline object created from a MATLAB expression can be used. Maximization is achieved by supplying the routines with `-f`, where `f` is the function being optimized.

Optimization options passed to the routines change optimization parameters. Default optimization parameters are used extensively but can be changed through an options structure.

Gradients are calculated using an adaptive finite-difference method unless they are supplied in a function. Parameters can be passed directly to functions, avoiding the need for global variables.

This *User's Guide* separates “medium-scale” algorithms from “large-scale” algorithms. Medium-scale is not a standard term and is used here only to differentiate these algorithms from the large-scale algorithms, which are designed to handle large-scale problems efficiently.

Medium-Scale Algorithms

The Optimization Toolbox routines offer a choice of algorithms and line search strategies. The principal algorithms for unconstrained minimization are the Nelder-Mead simplex search method and the BFGS (Broyden, Fletcher, Goldfarb, and Shanno) quasi-Newton method. For constrained minimization, minimax, goal attainment, and semi-infinite optimization, variations of *sequential quadratic programming* (SQP) are used. Nonlinear least squares problems use the Gauss-Newton and Levenberg-Marquardt methods.

A choice of line search strategy is given for unconstrained minimization and nonlinear least squares problems. The line search strategies use safeguarded cubic and quadratic interpolation and extrapolation methods.

Large-Scale Algorithms

All the large-scale algorithms, except linear programming, are trust-region methods. Bound constrained problems are solved using reflective Newton methods. Equality constrained problems are solved using a projective preconditioned conjugate gradient iteration. You can use sparse iterative solvers or sparse direct solvers in solving the linear systems to determine the current step. Some choice of preconditioning in the iterative solvers is also available.

The linear programming method is a variant of Mehrotra's predictor-corrector algorithm, a primal-dual interior-point method.

Examples that Use Standard Algorithms

This section presents the *medium-scale* (i.e., standard) algorithms through a tutorial. Examples similar to those in the first part of this tutorial (through the “Equality Constrained Example”) can also be found in the first demonstration “Tutorial Walk Through” in the M-file `optdemo`. The examples in this manual differ in that they use M-file functions for the objective functions, whereas the online demonstrations use inline objects for some functions.

Note Medium-scale is not a standard term and is used to differentiate these algorithms from the large-scale algorithms described in the “Large-Scale Algorithms” chapter.

The tutorial discusses the functions `fminunc` and `fmincon` in detail. The other optimization routines `fgoalattain`, `fminimax`, `lsqnonlin`, `fsolve`, and `fseminf` are used in a nearly identical manner, with differences only in the problem formulation and the termination criteria. The section “Multiobjective Examples” on page 1-18 discusses multiobjective optimization and gives several examples using `lsqnonlin`, `fminimax`, and `fgoalattain` including how Simulink can be used in conjunction with the toolbox.

This section includes the following examples:

- Unconstrained Example
- Nonlinear Inequality Constrained Example
- Constrained Example with Bounds
- Constrained Example with Gradients
- Gradient Check: Analytic Versus Numeric
- Equality Constrained Example

It also discusses:

- Maximization
- Greater-Than-Zero Constraints
- Additional Arguments: Avoiding Global Variables
- Multiobjective Examples

Unconstrained Example

Consider the problem of finding a set of values $[x_1, x_2]$ that solves

$$\underset{x}{\text{minimize}} \quad f(x) = e^{x_1}(4x_1^2 + 2x_2^2 + 4x_1x_2 + 2x_2 + 1) \quad (1-1)$$

To solve this two-dimensional problem, write an M-file that returns the function value. Then, invoke the unconstrained minimization routine `fminunc`.

Step 1: Write an M-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: Invoke one of the unconstrained optimization routines

```
x0 = [-1,1];    % Starting guess
options = optimset('LargeScale','off');
[x,fval,exitflag,output] = fminunc(@objfun,x0,options);
```

After 40 function evaluations, this produces the solution

```
x =
    0.5000   -1.0000
```

The function at the solution x is returned in `fval`.

```
fval =
    1.3030e-10
```

The `exitflag` tells if the algorithm converged. An `exitflag > 0` means a local minimum was found.

```
exitflag =
    1
```

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`, and the type of algorithm used in `algorithm`.

```
output =
    iterations: 7
    funcCount: 40
    stepsize: 1
    firstorderopt: 9.2801e-004
    algorithm: 'medium-scale: Quasi-Newton line search'
```

When more than one local minimum exists, the initial guess for the vector $[x_1, x_2]$ affects both the number of function evaluations and the value of the solution point. In the example above, x_0 is initialized to $[-1, 1]$.

The variable options can be passed to `fminunc` to change characteristics of the optimization algorithm, as in

```
x = fminunc(@objfun,x0,options);
```

`options` is a structure that contains values for termination tolerances and algorithm choices. An `options` structure can be created using the `optimset` function

```
options = optimset('LargeScale','off');
```

In this example we have turned off the default selection of the large-scale algorithm and so the medium-scale algorithm is used. Other options include controlling the amount of command line display during the optimization iteration, the tolerances for the termination criteria, if a user-supplied gradient or Jacobian is to be used, and the maximum number of iterations or function evaluations. See `optimset`, the individual optimization functions, and Table 4-3, Optimization Options Parameters, on page 4-11 for more options and information.

Nonlinear Inequality Constrained Example

If inequality constraints are added to Eq. 1-1, the resulting problem may be solved by the `fmincon` function. For example, find x that solves

$$\underset{x}{\text{minimize}} \ f(x) = e^{x_1}(4x_1^2 + 2x_2^2 + 4x_1x_2 + 2x_2 + 1) \quad (1-2)$$

subject to the constraints

$$x_1 x_2 - x_1 - x_2 \leq -1.5$$

$$x_1 x_2 \geq -10$$

Since neither of the constraints is linear, you cannot pass the constraints to `fmincon` at the command line. Instead you can create a second M-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

$$\begin{aligned} x_1 x_2 - x_1 - x_2 + 1.5 &\leq 0 \\ -x_1 x_2 - 10 &\leq 0 \end{aligned} \tag{1-3}$$

Step 1: Write an M-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 2: Invoke constrained optimization routine

```
x0 = [-1,1]; % Make a starting guess at the solution
options = optimset('LargeScale','off');
[x, fval] = ...
fmincon(@objfun,x0,[],[],[],[],[],[],@confun,options)
```

After 38 function calls, the solution `x` produced with function value `fval` is

```
x =
   -9.5474    1.0474
fval =
    0.0236
```

We can evaluate the constraints at the solution

```
[c,ceq] = confun(x)
```

```
c =
    1.0e-14 *
    0.1110
   -0.1776

ceq =
    []
```

Note that both constraint values are less than or equal to zero, that is, x satisfies $c(x) \leq 0$.

Constrained Example with Bounds

The variables in x can be restricted to certain limits by specifying simple bound constraints to the constrained optimizer function. For `fmincon`, the command

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

limits x to be within the range $lb \leq x \leq ub$.

To restrict x in Eq. 1-2 to be greater than zero (i.e., $x_1 \geq 0$, $x_2 \geq 0$), use the commands

```
x0 = [-1,1];           % Make a starting guess at the solution
lb = [0,0];           % Set lower bounds
ub = [ ];             % No upper bounds
options = optimset('LargeScale','off');
[x,fval] = ...
    fmincon(@objfun,x0,[],[],[],[],lb,ub,@confun,options)
[c, ceq] = confun(x)
```

Note that to pass in the lower bounds as the seventh argument to `fmincon`, you must specify values for the third through sixth arguments. In this example, we specified `[]` for these arguments since we have no linear inequalities or linear equalities.

After 13 function evaluations, the solution produced is

```
x =
    0    1.5000
fval =
    8.5000
```



```

c =
    0
   -10
ceq =
    []

```

When `lb` or `ub` contains fewer elements than `x`, only the first corresponding elements in `x` are bounded. Alternatively, if only some of the variables are bounded, then use `-inf` in `lb` for unbounded below variables and `inf` in `ub` for unbounded above variables. For example,

```

lb = [-inf 0];
ub = [10 inf];

```

bounds $x_1 \leq 10$, $0 \leq x_2$ (x_1 has no lower bound and x_2 has no upper bound). Using `inf` and `-inf` give better numerical results than using a very large positive number or a very large negative number to imply lack of bounds.

Note that the number of function evaluations to find the solution is reduced since we further restricted the search space. Fewer function evaluations are usually taken when a problem has more constraints and bound limitations because the optimization makes better decisions regarding step-size and regions of feasibility than in the unconstrained case. It is, therefore, good practice to bound and constrain problems, where possible, to promote fast convergence to a solution.

Constrained Example 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 the Eq. 1-2 using analytically determined gradients, do the following:

Step 1: Write an M-file for the objective function and gradient

```

function [f,G] = 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

```

```
t = exp(x(1))*(4*x(1)^2+2*x(2)^2+4*x(1)*x(2)+2*x(2)+1);
G = [ t + exp(x(1)) * (8*x(1) + 4*x(2)),
      exp(x(1))*(4*x(1)+4*x(2)+2)];
```

Step 2: Write an M-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;
% Gradient of the constraints
DC= [x(2)-1, -x(2);
      x(1)-1, -x(1)];
% No nonlinear equality constraints
ceq=[];
DCEq = [ ];
```

G contains the partial derivatives of the objective function, **f**, returned by `objfungrad(x)`, with respect to each of the elements in **x**.

$$\frac{\partial f}{\partial x} = \begin{bmatrix} e^{x_1}(4x_1^2 + 2x_2^2 + 4x_1x_2 + 2x_2 + 1) + e^{x_1}(8x_1 + 4x_2) \\ e^{x_1}(4x_1 + 4x_2 + 2) \end{bmatrix} \quad (1-4)$$

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

$$\begin{bmatrix} \frac{\partial c_1}{\partial x_1} & \frac{\partial c_2}{\partial x_1} \\ \frac{\partial c_1}{\partial x_2} & \frac{\partial c_2}{\partial x_2} \end{bmatrix} = \begin{bmatrix} x_2 - 1 & -x_2 \\ x_1 - 1 & -x_1 \end{bmatrix} \quad (1-5)$$

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 M-files contain this additional information. Use `optimset` to turn the

parameters `GradObj` and `GradConstr` to 'on' in the example's existing options structure

```
options = optimset(options, 'GradObj', 'on', 'GradConstr', 'on');
```

If you do not set these parameters to 'on' in the options structure, `fmincon` will not use the analytic gradients.

The arguments `lb` and `ub` place lower and upper bounds on the independent variables in `x`. In this example we have no bound constraints and so they are both set to `[]`.

Step 3: Invoke constrained optimization routine

```
x0 = [-1,1]; % Starting guess
options = optimset('LargeScale','off');
options = optimset(options, 'GradObj', 'on', 'GradConstr', 'on');
lb = [ ]; ub = [ ]; % No upper or lower bounds
[x,fval] = fmincon(@objfungrad,x0,[],[],[],[],lb,ub,...
    @confungrad,options)
[c,ceq] = confungrad(x) % Check the constraint values at x
```

After 20 function evaluations, the solution produced is

```
x =
    -9.5474    1.0474
fval =
    0.0236
c =
    1.0e-14 *
    0.1110
    -0.1776
ceq =
    []
```

Gradient Check: Analytic Versus Numeric

When analytically determined gradients are provided, you can compare the supplied gradients with a set calculated by finite-difference evaluation. This is particularly useful for detecting mistakes in either the objective function or the gradient function formulation.

If such gradient checks are desired, set the `DerivativeCheck` parameter to 'on' using `optimset`.

```
options = optimset(options, 'DerivativeCheck', 'on');
```

The first cycle of the optimization checks the analytically determined gradients (of the objective function and, if they exist, the nonlinear constraints). If they do not match the finite-differencing gradients within a given tolerance, a warning message indicates the discrepancy and gives the option to abort the optimization or to continue.

Equality Constrained Example

For routines that permit equality constraints, nonlinear equality constraints must be computed in the M-file 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 $x_1 x_2 \geq -10$, rewrite them as

$$\begin{aligned}x_1^2 + x_2 - 1 &= 0 \\ -x_1 x_2 - 10 &\leq 0\end{aligned}$$

and then solve the problem using the following steps:

Step 1: Write an M-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 an M-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 = optimset('LargeScale','off');
[x,fval] = fmincon(@objfun,x0,[],[],[],[],[],[],...
    @confuneq,options)
[c,ceq] = confuneq(x) % Check the constraint values at x

```

After 21 function evaluations, the solution produced is

```

x =
    -0.7529    0.4332
fval =
    1.5093
c =
   -9.6739
ceq =
    4.0684e-010

```

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 zero as desired.

Maximization

The optimization functions `fminbnd`, `fminsearch`, `fminunc`, `fmincon`, `fgoalattain`, `fminimax`, `lsqcurvefit`, and `lsqnonlin` all perform minimization of the objective function, $f(x)$. Maximization is achieved by supplying the routines with $-f(x)$. Similarly, to achieve maximization for `quadprog` supply `-H` and `-f`, and for `linprog` supply `-f`.

Greater-Than-Zero Constraints

The Optimization Toolbox assumes nonlinear inequality constraints are of the form $C_i(x) \leq 0$. Greater-than-zero constraints are expressed as less than zero 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 $C_i(x) \geq b$ is equivalent to the constraint $-C_i(x) + b \leq 0$.

Additional Arguments: Avoiding Global Variables

Parameters that would otherwise have to be declared as global can be passed directly to M-file functions using additional arguments at the end of the calling sequence.

For example, entering a number of variables at the end of the call to `fsolve`

```
[x,fval] = fsolve(@objfun,x0,options,P1,P2,...)
```

passes the arguments directly to the functions `objfun` when it is called from inside `fsolve`.

```
F = objfun(x,P1,P2, ... )
```

Consider, for example, finding zeros of the function `ellipj(u,m)`. The function needs parameter `m` as well as input `u`. To look for a zero near `u = 3`, for `m = 0.5`

```
m = 0.5;
options = optimset('Display','off'); % Turn off Display
x = fsolve(@ellipj,3,options,m)
```

returns

```
x =
    3.7081
```

Then, solve for the function `ellipj`.

```
f = ellipj(x,m)
f =
   -2.9925e-008
```

The call to `optimset` to get the default options for `fsolve` implies that default tolerances are used and that analytic gradients are not provided.

Multiobjective Examples

The previous examples involved problems with a single objective function. This section demonstrates solving problems with multiobjective functions using `lsqnonlin`, `fminimax`, and `fgoalattain`. Included is an example of how to optimize parameters in a Simulink model.

Simulink Example

Let's say that you want to optimize the control parameters in the Simulink model `optsim.mdl`. (This model can be found in the Optimization Toolbox `optim` directory. 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 shown in Figure 1-1, Plant with Actuator Saturation.




Figure 1-1: 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 open-loop response of the system to a step input is shown in Figure 1-2, Open-Loop Response. You can see this response by opening the model (type `optsim` at the command line, or click the model name), opening the Scope block, and running the simulation. The response plots to the Scope.




Figure 1-2: Open-Loop Response

The problem is to design a feedback control law 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. See Figure 1-3, Closed-Loop Model.

Tunable Variables are PID gains, K_p , K_i , and K_d .




Figure 1-3: 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 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. This is defined via a MATLAB function in the file `tracklsq.m` shown below that defines the error signal. The error signal is `yout`, the output computed by calling `sim`, minus the input signal `1`.

The function `tracklsq` must run the simulation. The simulation can be run either in the base workspace or the current workspace, i.e., the workspace of the function calling `sim`, which in this case is `tracklsq`'s workspace. In this example, the `simset` command is used to tell `sim` to run the simulation in the current workspace by setting `'SrcWorkspace'` to `'Current'`.

To run the simulation in `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 we are optimizing. You can initialize `a1` and `a2` before calling `lsqnonlin` and then pass these two variables as additional arguments. `lsqnonlin` then passes `a1` and `a2` to `tracklsq` each time it is called, so you do not have to use global variables.

After choosing a solver using the `simset` function, the simulation is run using `sim`. The simulation is performed using a fixed-step fifth-order method to 100 seconds. When the simulation completes, the variables `tout`, `xout`, and `yout` are now in the current workspace (that is, the workspace of `tracklsq`). The Output block is used in the block diagram model to put `yout` into the current workspace at the end of the simulation.

Step 1: Write an M-file `tracklsq.m`

```
function F = tracklsq(pid,a1,a2)
Kp = pid(1);          % Move variables into model parameter names
Ki = pid(2);
Kd = pid(3);
% Choose solver and set model workspace to this function
opt = simset('solver','ode5','SrcWorkspace','Current');
[tout,xout,yout] = sim('optsim',[0 100],opt);
F = yout-1;          % Compute error signal
```

Step 2: Invoke optimization routine

```

optsim                                % Load the model
pid0 = [0.63 0.0504 1.9688] % Set initial values
a1 = 3; a2 = 43;                    % Initialize plant variables in model
options = optimset('LargeScale','off','Display','iter',...
                  'TolX',0.001,'TolFun',0.001);
pid = lsqnonlin(@tracklsq, pid0, [], [], options, a1, a2)
% Put variables back in the base workspace
Kp = pid(1); Ki = pid(2); Kd = pid(3);

```

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.

The optimization gives the solution for the proportional, integral, and derivative (Kp, Ki, Kd) gains of the controller after 73 function evaluations.

Iteration	Func-count	Residual	Step-size	Directional derivative	Lambda
1	3	8.66531	1	-3.48	
2	10	6.78831	1	-0.0634	3.4355
3	19	5.99204	5.5	-0.0446	0.28612
4	28	4.74992	5.78	-0.0213	0.0227966
5	36	4.51795	1.25	0.0222	0.0744258
6	43	4.5115	0.58	-0.00633	0.03445
7	51	4.49455	2.99	0.000688	0.017225
8	58	4.4836	0.915	0.00203	0.0180998
9	66	4.47724	1.22	0.000845	0.00904992
10	73	4.47405	0.801	-0.00072	0.0113409

```

Optimization terminated successfully:
  Gradient in the search direction less than tolFun
  Gradient less than 10*(tolFun+tolX)
pid =
    2.9108    0.1443   12.8161

```

The resulting closed-loop step response is shown in Figure 1-4.




Figure 1-4: 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 may 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 may 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 [1].)

The Nonlinear Control Design (NCD) Blockset is recommended for solving multiobjective optimization problems in conjunction with variable-step solvers in Simulink. It provides a special numeric gradient computation that works with Simulink and avoids introducing the lack of smoothness problem.

Another solution approach 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. Then in the function `trackmmobj` the objective function is simply the output `yout` returned by the `sim` command. But minimizing the maximum output at all time steps may force the output far below unity for some time steps. To keep the output above 0.95 after the first 20 seconds, in the constraint function `trackkmmcon` add a constraint `yout >= 0.95` from $t=20$ to $t=100$. Because constraints must be in the form `g <= 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. The nonlinear constraint function is *always* called immediately after the objective function in `fmincon`, `fminimax`, `fgoalattain`, and `fseminf` with the same values. Thus you can avoid calling the simulation twice by using `assignin` to assign the current value of `F` to the variable `F_TRACKMMOBJ` in the base workspace. Then the first step in `trackmmcon` is to use `evalin` to evaluate the variable `F_TRACKMMOBJ` in the base workspace, and assign the result to `F` locally in `trackmmcon`.

Step 1: Write an M-file trackmmobj.m to compute objective function

```
function F = trackmmobj(pid,a1,a2)
Kp = pid(1);
Ki = pid(2);
Kd = pid(3);
% Compute function value
opt = simset('solver','ode5','SrcWorkspace','Current');
[tout,xout,yout] = sim('optsim',[0 100],opt);
F = yout;
assignin('base','F_TRACKMMOBJ',F);
```

Step 2: Write an M-file trackmmcon.m to compute nonlinear constraints

```
function [c,ceq] = trackmmcon(pid,a1,a2)
F = evalin('base','F_TRACKMMOBJ');
% Compute constraints
c = -F(20:100)+.95;
ceq = [ ];
```

Note that `fminimax` passes `a1` and `a2` to the objective and constraint values, so `trackmmcon` needs input arguments for these variables even though it does not use them.

Step 3: Invoke constrained optimization routine

```
optsim
pid0 = [0.63 0.0504 1.9688]
a1 = 3; a2 = 43;
options = optimset('Display','iter',...
                  'TolX',0.001,'TolFun',0.001);
pid = fminimax(@trackmmobj,pid0,[],[],[],[],[],[],...
              'trackmmcon',options,a1,a2)
% Put variables back in the base workspace
Kp = pid(1); Ki = pid(2); Kd = pid(3);
```

resulting in

Iter	F-count	Max {F,constraints}	Step-size	Directional derivative	Procedure
1	5	1.12	1	1.18	
2	11	1.264	1	-0.172	

3	17	1.055	1	-0.0128	Hessian modified twice
4	23	1.004	1	3.48e-005	Hessian modified
5	29	0.9997	1	-1.36e-006	Hessian modified twice

Optimization terminated successfully:
Search direction less than 2*options.TolX and
maximum constraint violation is less than options.TolCon

Active Constraints:

```
1
14
182
pid =
0.5894    0.0605    5.5295
```

The last value shown in the $\text{MAX}\{F, \text{constraints}\}$ column of the output shows the maximum value for all the time steps is 0.9997. The closed loop response with this result is shown in Figure 1-5, Closed-Loop Response Using `fminimax`.

This solution differs from the `lsqnonlin` solution as you are solving different problem formulations.




Figure 1-5: Closed-Loop Response Using `fminimax`

Signal Processing Example

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

$$\begin{aligned}
 H(f) &= \sum_{n=0}^{2M} h(n) e^{-j2\pi fn} \\
 &= A(f) e^{-j2\pi fM} \\
 A(f) &= \sum_{n=0}^{M-1} a(n) \cos(2\pi fn)
 \end{aligned}
 \tag{1-6}$$

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, the function `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 takes `a`, the magnitude function coefficients, and `w`, the discretization of the frequency domain we are interested in.

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 2-42). 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 = optimset('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 an M-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.'-:');
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 = optimset('GoalsExactAchieve',length(goal));
[a,fval,attainfactor,exitflag] = fgoalattain(@filtmin,...
    a0,goal,weight,[],[],[],[],[],[],[],options,w0);

% 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 (Figure 1-6). Note that you could use the `remez` function in the Signal Processing Toolbox to design this filter.




Figure 1-6: Magnitude Response with Initial and Final Magnitude Coefficients

Large-Scale Examples

Some of the optimization functions include algorithms for continuous optimization problems especially targeted to large problems with sparsity or structure. The main large-scale algorithms are iterative, i.e., a sequence of approximate solutions is generated. In each iteration a linear system is (approximately) solved. The linear systems are solved using the sparse matrix capabilities of MATLAB and a variety of sparse linear solution techniques, both iterative and direct.

Generally speaking the large-scale optimization methods preserve structure and sparsity, using exact derivative information wherever possible. To solve the large-scale problems efficiently, some problem formulations are restricted (such as only solving overdetermined linear or nonlinear systems), or require additional information (e.g., the nonlinear minimization algorithm requires the gradient be computed in the user-supplied function).

This section summarizes the kinds of problems covered by large-scale methods and provides these examples:

- Nonlinear Equations with Jacobian
- Nonlinear Equations with Jacobian Sparsity Pattern
- Nonlinear Least Squares with Full Jacobian Sparsity Pattern
- Nonlinear Minimization with Gradient and Hessian
- Nonlinear Minimization with Gradient and Hessian Sparsity Pattern
- Nonlinear Minimization with Bound Constraints and Banded Preconditioner
- Nonlinear Minimization with Equality Constraints
- Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints
- Quadratic Minimization with Bound Constraints
- Quadratic Minimization with a Dense but Structured Hessian
- Linear Least Squares with Bound Constraints
- Linear Programming with Equalities and Inequalities
- Linear Programming with Dense Columns in the Equalities

Problems Covered by Large-Scale Methods

Not all possible problem formulations are covered by the large-scale algorithms. The following table describes what functionality is covered by the large-scale methods. For example, for `fmincon`, the large-scale algorithm covers the case where there are only bound constraints or only linear equalities. For each problem formulation, the third column describes what additional information is needed for the large-scale algorithms. For `fminunc` and `fmincon`, the gradient must be computed along with the objective in the user-supplied function (the gradient is not required for the medium-scale algorithms).

Since these methods can also be used on small- to medium-scale problems that are not necessarily sparse, the last column of the table emphasizes what conditions are needed for large-scale problems to run efficiently without exceeding your computer system's memory capabilities, e.g., the linear constraint matrices should be sparse. For smaller problems the conditions in the last column are unnecessary.

Note The following tables lists the functions in order of increasing problem complexity.

Several examples, which follow this table, clarify the contents of the table.

Table 1-4: Large-Scale Problem Coverage and Requirements

Function	Problem Formulations	Additional Information Needed	For Large Problems
<code>fminunc</code>	$\min_x f(x)$	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.

Table 1-4: Large-Scale Problem Coverage and Requirements (Continued)

Function	Problem Formulations	Additional Information Needed	For Large Problems
fmincon	<ul style="list-style-type: none"> • $\min_x f(x)$ such that $l \leq x \leq u$ where $l < u$ • $\min_x f(x)$ such that $Aeq \cdot x = beq$, and Aeq is an m-by-n matrix where $m \leq n$. 	Must provide gradient for $f(x)$ in fun.	<ul style="list-style-type: none"> • Provide sparsity structure of the Hessian, or compute the Hessian in fun. • The Hessian should be sparse. • Aeq should be sparse.
lsqnonlin	<ul style="list-style-type: none"> • $\min_x \frac{1}{2} \ F(x)\ _2^2 = \frac{1}{2} \sum_i F_i(x)^2$ • $\min_x \frac{1}{2} \ F(x)\ _2^2 = \frac{1}{2} \sum_i F_i(x)^2$ such that $l \leq x \leq u$ where $l < u$ <p>$F(x)$ must be overdetermined (have at least as many equations as variables).</p>	None	<ul style="list-style-type: none"> • Provide sparsity structure of the Jacobian, or compute the Jacobian in fun. • The Jacobian should be sparse.
lsqcurvefit	<ul style="list-style-type: none"> • $\min_x \frac{1}{2} \ F(x, xdata) - ydata\ _2^2$ • $\min_x \frac{1}{2} \ F(x, xdata) - ydata\ _2^2$ such that $l \leq x \leq u$ where $l < u$ <p>$F(x, xdata)$ must be overdetermined (have at least as many equations as variables).</p>	None	<ul style="list-style-type: none"> • Provide sparsity structure of the Jacobian, or compute the Jacobian in fun. • The Jacobian should be sparse.

Table 1-4: Large-Scale Problem Coverage and Requirements (Continued)

Function	Problem Formulations	Additional Information Needed	For Large Problems
fsolve	$F(x) = 0$ $F(x)$ must have the same number of equations as variables.	None	<ul style="list-style-type: none"> • Provide sparsity structure of the Jacobian or compute the Jacobian in fun. • The Jacobian should be sparse.
lsqlin	$\min_x \ C \cdot x - d\ _2^2$ such that $l \leq x \leq u$ where $l < u$ C is an m-by-n matrix where $m \geq n$, i.e., the problem must be overdetermined.	None	C should be sparse.
linprog	$\min_x f^T x$ such that $A \cdot x \leq b$ and $Aeq \cdot x = beq$, where $l \leq x \leq u$	None	A and Aeq should be sparse.
quadprog	<ul style="list-style-type: none"> • $\min_x \frac{1}{2} x^T H x + f^T x$ such that $l \leq x \leq u$ where $l < u$ • $\min_x \frac{1}{2} x^T H x + f^T x$ such that $Aeq \cdot x = beq$, and Aeq is an m-by-n matrix where $m \leq n$. 	None	<ul style="list-style-type: none"> • H should be sparse. • Aeq should be sparse.

In the examples below, many of the M-file functions are available in the Optimization Toolbox `optim` directory. Most of these do not have a fixed problem size, i.e., the size of your starting point `xstart` will determine the size problem that is computed. If your computer system cannot handle the size suggested in the examples below, use a smaller dimension start point to run the problems. If the problems have upper or lower bounds or equalities, you will have to adjust the size of those vectors or matrices as well.

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, $F(x) = 0$, use the large-scale method available in `fsolve`.

Step 1: Write an M-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 = optimset('Display','iter','Jacobian','on');
[x,fval,exitflag,output] = fsolve(fun,xstart,options);

```

A starting point is given as well as the function name. The large-scale method is the default, so it is not necessary to specify this in the options argument. Output at each iteration is specified. Finally, so that `fsolve` will use the Jacobian information available in `nlsf1.m`, you need to turn the Jacobian parameter 'on' using `optimset`.

The commands display this output.

Iteration	Func-count	f(x)	Norm of step	First-order optimality	CG-Iterations
1	2	1011	1	19	0
2	3	16.1942	7.91898	2.35	3
3	4	0.0228027	1.33142	0.291	3
4	5	0.000103359	0.0433329	0.0201	4
5	6	7.3792e-007	0.0022606	0.000946	4
6	7	4.02299e-010	0.000268381	4.12e-005	5

Optimization terminated successfully:
Relative function value changing by less than `OPTIONS.TolFun`

A linear system is (approximately) solved in each major iteration using the preconditioned conjugate gradient method. The default value for `PrecondBandWidth` is 0 in options, so 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 5 for a problem of 1000 dimensions, implying the linear systems are not very difficult to solve in this case (though more work is required as convergence progresses).

It is possible to override the default choice of preconditioner (diagonal) by choosing a banded preconditioner through the use of the parameter

PrecondBandWidth. 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 = optimset('Display','iter','Jacobian','on',...
                  'PrecondBandWidth',1) ;
[x,fval,exitflag,output] = fsolve(fun,xstart,options) ;
```

In this case the output is

Iteration	Func-count	f(x)	Norm of First-order step	optimality	CG- Iterations
1	2	1011	1	19	0
2	3	16.0839	7.92496	1.92	1
3	4	0.0458181	1.3279	0.579	1
4	5	0.000101184	0.0631898	0.0203	2
5	6	3.16615e-007	0.00273698	0.00079	2
6	7	9.72481e-010	0.00018111	5.82e-005	2

Optimization terminated successfully:

Relative function value changing by less than OPTIONS.TolFun

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 Gradients” on page 3-6.

Nonlinear Equations with Jacobian Sparsity Pattern

In the preceding example 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 the Jacobian can be computed in `nlsf1` (using the Jacobian parameter in options), `fsolve`, `lsqnonlin`, and `lsqcurvefit` will instead use finite-differencing to approximate the Jacobian.

In order for this finite-differencing to be as efficient as possible, the sparsity pattern of the Jacobian should be supplied, using the `JacobPattern` parameter 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 `nlsfs1`, the finite-differencing code will attempt 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 will attempt to compute. In other words, this problem is solvable if the sparsity pattern is provided. If not, most computers will 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 the same as `nlsf1` except only the function values are returned; sparse finite-differencing is used to estimate the sparse Jacobian matrix as needed.

Step 1: Write an M-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 = optimset('Display','iter','JacobPattern',Jstr,...
                  'PrecondBandWidth',1);
[x,fval,exitflag,output] = fsolve(fun,xstart,options);
```

In this case, the output displayed is

Iteration	Func-count	f(x)	Norm of First-order step	First-order optimality	CG- Iterations
1	6	1011	1	19	0
2	11	16.0839	7.92496	1.92	1
3	16	0.0458181	1.3279	0.579	1

```

4      21      0.000101184    0.0631898    0.0203      2
5      26      3.16615e-007    0.00273698    0.00079      2
6      31      9.72482e-010    0.00018111    5.82e-005     2

```

Optimization terminated successfully:
Relative function value changing by less than OPTIONS.TolFun

Alternatively, it is possible to choose a sparse direct linear solver (i.e., a sparse QR factorization) by indicating a “complete” preconditioner, i.e., if we set PrecondBandWidth to Inf, then a sparse direct linear solver will be used instead of a preconditioned conjugate gradient iteration

```

xstart = -ones(1000,1);
fun = @nlsf1a;
load nlsdat1 % Get Jstr
options = optimset('Display','iter','JacobPattern',Jstr,...
                  'PrecondBandWidth',inf);
[x,fval,exitflag,output] = fsolve(fun,xstart,options);

```

and the resulting display is

Iteration	Func-count	f(x)	Norm of First-order step	CG-optimality	CG-Iterations
1	6	1011	1	19	0
2	11	15.9018	7.92421	1.89	1
3	16	0.0128163	1.32542	0.0746	1
4	21	1.73538e-008	0.0397925	0.000196	1
5	26	1.13169e-018	4.55544e-005	2.76e-009	1

Optimization terminated successfully:
Relative function value changing by less than OPTIONS.TolFun

When the sparse direct solvers are used, the CG iteration will be 1 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 Least Squares with Full Jacobian Sparsity Pattern

The large-scale methods in `lsqnonlin`, `lsqcurvefit`, and `fsolve` can be used with small- to medium-scale problems without computing the Jacobian in `fun` or providing the Jacobian sparsity pattern. (This example also applies to the case of 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 available 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 may still be too large, but you can only find out by running it and seeing if MATLAB is able to run within the amount of virtual memory available on your system.

Let's say you have a small problem with 10 equations and 2 unknowns, such as 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 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 k components).

Step 1: Write an M-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] = lsqnonlin(@myfun,x0) % Invoke optimizer
```

Since the Jacobian is not computed in `myfun.m`, and no Jacobian sparsity pattern is provided using the `JacobPattern` parameter in options, `lsqnonlin` calls the large-scale method with `JacobPattern` set to

`Jstr = sparse(ones(10,2))`. This is the default for `lsqnonlin`. Note that the Jacobian parameter in options is 'off' by default.

When the finite-differencing routine is called the first time, it detects that `Jstr` is actually a dense matrix, i.e., that 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 =
    0.2578    0.2578
resnorm      % Residual or sum of squares
resnorm =
    124.3622
```

Most computer systems will be able to handle much larger full problems, say into the 100's 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 will always run faster if this information is provided.

Nonlinear 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

$$f(x) = \sum_{i=1}^{n-1} \left[(x_i^2)^{(x_{i+1}^2+1)} + (x_{i+1}^2)^{(x_i^2+1)} \right] \quad (1-7)$$

where $n = 1000$.

Step 1: Write an M-file `brownfgh.m` that computes the objective function, the gradient of the objective, and the sparse tridiagonal Hessian matrix

This file is rather lengthy and is not included here. You can 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 `optimset` to indicate this information is available in `brownfgh` using the `GradObj` and `Hessian` parameters.

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 = optimset('GradObj',...
                  'on','Hessian','on');
[x,fval,exitflag,output] = fminunc(@brownfgh,xstart,options);
```

This 1000 variable problem is solved in 8 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.

```
exitflag =
     1
fval =
  2.8709e-017
output.iterations
ans =
     8
output.cgiterations
ans =
     7
output.firstorderopt
ans =
  4.7948e-010
```

Nonlinear Minimization with Gradient and Hessian Sparsity Pattern

Next we solve the same problem but the Hessian matrix is now approximated by sparse finite-differences instead of explicit computation. To use the large-scale method in `fminunc`, you *must* compute the gradient in `fun`; it is *not optional* as in the medium-scale method.

The M-file function `brownfg` computes the objective function and gradient.

Step 1: Write an M-file `brownfg.m` that computes the objective function and the gradient of the objective

```
function [f,g] = brownfg(x,dummy)
% 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 `optimset` to set the `HessPattern` parameter to `Hstr`. When a problem as large as this has obvious sparsity structure, not setting the `HessPattern` parameter will require a huge amount of unnecessary memory and computation since `fminunc` will attempt to use finite-differencing on a full Hessian matrix of one million nonzero entries.

You must also set the `GradObj` parameter to 'on' using `optimset` 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 = optimset('GradObj','on','HessPattern',Hstr);
[x,fval,exitflag,output] = fminunc(fun,xstart,options);

```

This 1000 variable problem is solved in 8 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).

```

exitflag =
    1
fval =
    7.4738e-017
output.iterations
ans =
    8
output.cgiterations
ans =
    7
output.firstorderopt
ans =
    7.9822e-010

```

Nonlinear Minimization with Bound Constraints and Banded Preconditioner

The goal in this problem is to minimize the nonlinear function

$$f(x) = 1 + \sum_{i=1}^n |(3 - 2x_i)x_i - x_{i-1} - x_{i+1} + 1|^p + \sum_{i=1}^{\frac{n}{2}} |x_i + x_{i+n/2}|^p$$

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 an M-file `tbroyfg.m` that computes the objective function and the gradient of the objective

The M-file function, `tbroyfg.m`, computes the function value and gradient. This file is lengthy 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 spy plot below.

```
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 `optimset` 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` will attempt to use finite-differencing on a full Hessian matrix of 640,000 nonzero entries.

You must also set the `GradObj` parameter to `'on'` using `optimset` 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 = optimset('GradObj','on','HessPattern',Hstr);
[x,fval,exitflag,output] = ...
    fmincon(fun,xstart,[],[],[],[],lb,ub,[],options);
```

After eight iterations, the `exitflag`, `fval` and output values are

```
exitflag =
    1
fval =
    270.4790
output =
    iterations: 8
    funcCount: 8
    cgiterations: 18
    firstorderopt: 0.0163
    algorithm: 'large-scale: trust-region reflective Newton'
```

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 3-8, 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 `optimset` 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 = optimset('GradObj','on','HessPattern',Hstr, ...
    '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 =
    1
fval =
    2.7048e+002
```

```

output =
    iterations: 10
    funcCount: 10
    cgiterations: 15
    firstorderopt: 7.5339e-005
    algorithm: 'large-scale: trust-region reflective Newton'

```

Nonlinear Minimization with Equality Constraints

The large-scale method for `fmincon` can handle equality constraints if no other constraints exist. Suppose you want to minimize the same objective in Eq. 1-7, which is coded in the function `brownfgh.m`, where $n = 1000$, such that $A_{eq} \cdot x = b_{eq}$ for A_{eq} that has 100 equations (so A_{eq} is a 100-by-1000 matrix).

Step 1: Write an M-file `brownfgh.m` that computes the objective function, the gradient of the objective, and the sparse tridiagonal Hessian matrix

As before, this file is rather lengthy and is not included here. You can 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 `optimset` to indicate this information is available in `brownfgh` using the `GradObj` and `Hessian` parameters.

The sparse matrix A_{eq} and vector b_{eq} 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.

```

condest(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 = optimset('GradObj','on','Hessian','on', ...
                  'PrecondBandWidth', inf);
[x,fval,exitflag,output] = ...
    fmincon(fun,xstart,[],[],Aeq,beq,[],[],[],options);
```

Setting the parameter `PrecondBandWidth` to `inf` causes a sparse direct solver to be used instead of preconditioned conjugate gradients.

The `exitflag` indicates convergence with the final function value `fval` after 16 iterations.

```
exitflag =
    1
fval =
    205.9313
output =
    iterations: 16
    funcCount: 16
    cgiterations: 14
    firstorderopt: 2.1434e-004
    algorithm: 'large-scale: projected trust-region Newton'
```

The linear equalities are satisfied at `x`.

```
norm(Aeq*x-beq)
ans =
    1.1913e-012
```

Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints

The `fmincon` and `fminunc` large-scale methods 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, as it does for medium-scale problems and for large-scale problems with sparse H , 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 objective function has the structure

$$f(x) = \hat{f}(x) - \frac{1}{2} x^T V V^T x$$

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

$$H = \hat{H} - V V^T$$

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

$$W = H*Y = (Hinfo - V*V')*Y$$

In this example, the Hessian multiply function needs \hat{H} and V to compute the Hessian matrix product. V is a constant and so V can be passed as an additional parameter to `fmincon`, then `fmincon` passes V as an additional parameter to `hmf1eq1`.

However \hat{H} is not a constant and must be computed at the current x . This can be done by computing \hat{H} in the objective function and returning \hat{H} , as `Hinfo`, in the third output argument. By using `optimset` 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 an M-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 lengthy 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 `optimset` to set the `GradObj` and Hessian parameters to `'on'`.

Step 2: Write a function to compute Hessian-matrix products for H given a matrix Y.

Now, define a function `hmfleq1` that uses `Hinfo`, which is computed in `brownvv`, and `V`, which the example passes to `fmincon` as an additional parameter, to compute the Hessian matrix product `W` where $W = H*Y = (Hinfo - V*V')*Y$. This function must have the form

```
W = hmfleq1(Hinfo,Y,p1,p2...)
```

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, any additional parameters passed to `fmincon` are passed to the Hessian multiply function, so `hmfleq1` must accept the same additional parameters, e.g. the matrix `V`.

```
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 Optimization Toolbox as M-file `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, A_{eq} and b_{eq} , from `fleq1.mat`, which is available in the Optimization Toolbox. Use `optimset` to set the `GradObj` and `Hessian` options '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.

```
load fleq1                % Get V, Aeq, beq
n = 1000;                 % problem dimension
mtxmpy = @hmfleq1;       % Function handle to function hmfleq1
xstart = -ones(n,1); xstart(2:2:n,1) = ones(length(2:2:n),1);
options = optimset('GradObj','on','Hessian','on',...
                  'HessMult',mtxmpy,'Display','iter');
[x,fval,exitflag,output] = fmincon(@brownvv,xstart,[],[],...
                                  Aeq,beq,[],[],[],...
                                  options,V);
```

Note Type `[fval,exitflag,output] = runfleq1` to run the code above. This command displays the values for `fval`, `exitflag`, and `output`, as well as the iterative display below.

Since the iterative display was set using `optimset`, the results printed to the command line window are

Iteration	f(x)	Norm of step	First-order optimality	CG-iterations
1	1997.07	1	555	0
2	1072.56	6.31716	377	1
3	480.232	8.19554	159	2
4	136.861	10.3015	59.5	2
5	44.3708	9.04697	16.3	2
6	44.3708	100	16.3	2
7	44.3708	25	16.3	0
8	-8.90967	6.25	28.5	0
9	-318.486	12.5	107	1
10	-318.486	12.5	107	1
11	-415.445	3.125	73.9	0

12	-561.688	3.125	47.4	2
13	-785.326	6.25	126	3
14	-785.326	4.30584	126	5
15	-804.414	1.07646	26.9	0
16	-822.399	2.16965	2.8	3
17	-823.173	0.40754	1.34	3
18	-823.241	0.154885	0.555	3
19	-823.246	0.0518407	0.214	5
2	-823.246	0.00977601	0.00724	6

Optimization terminated successfully:

Relative function value changing by less than OPTIONS.TolFun

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

```
norm(Aeq*x-beq) =
    1.2861e-013
```

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.

Quadratic Minimization with Bound Constraints

To minimize a large-scale quadratic with upper and lower bounds, you can use the `quadprog` function.

The problem stored in the MAT-file `qpbox1.mat` is a positive definite quadratic, and the Hessian matrix `H` is tridiagonal, subject to upper (`ub`) and lower (`lb`) bounds.

Load the Hessian and define f , lb , ub . Call a quadratic minimization routine with a starting point $xstart$

```
load qpbox1 % 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;
xstart = 0.5*ones(400,1);
[x,fval,exitflag,output] = ...
    quadprog(H,f,[],[],[],[],lb,ub,xstart);
```

Looking at the resulting values of $exitflag$ and $output$

```
exitflag =
    1
output =
    firstorderopt: 7.8435e-006
    iterations: 20
    cgiterations: 1809
    algorithm: 'large-scale: reflective trust-region'
```

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 = optimset('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 =
    1
output =
    firstorderopt: 2.3821e-005
    iterations: 36
    cgiterations: 1547
    algorithm: 'large-scale: reflective trust-region'
```

A second strategy would be to use a direct solver at each iteration by setting the `PrecondBandWidth` parameter to `inf`.

```
options = optimset('PrecondBandWidth',inf);
[x,fval,exitflag,output] = ...
    quadprog(H,f,[],[],[],[],lb,ub,xstart,options);
```

Now the number of iterations has dropped to 10.

```
exitflag =
    1
output =
    firstorderopt: 4.8955e-007
    iterations: 10
    cgiterations: 9
    algorithm: 'large-scale: reflective trust-region'
```

Using a direct solve at each iteration usually causes the number of iterations to decrease, but often takes more time per iteration. For this problem, the trade-off is beneficial as the time for `quadprog` to solve the problem decreases by a factor of 10.

Quadratic Minimization with a Dense but Structured Hessian

The `quadprog` large-scale 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 medium-scale problems and for large-scale 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 `qpbbox4mult`. 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 passed to the Hessian multiply function `qpbbox4mult` from `quadprog`. There are two ways to indicate this in the

call to `quadprog`. The first argument passed to `quadprog` is passed to the Hessian multiply function. Also, parameters passed to `quadprog` as additional parameters are passed to the Hessian multiply function.

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 1-57). A is then passed as an additional parameter.

```
quadprog(B,f,[],[],[],[],1,u,xstart,options,A)
```

Step 2: Write a function to compute Hessian-matrix products for H.

Now, define a function `qpbox4mult` that uses A and B to compute the Hessian matrix product W where $W = H*Y = (B + A*A')*Y$. This function must have the form

```
W = qpbox4mult(Hinfo,Y,p1,p2...)
```

`qpbox4mult` must accept the same first argument as passed to `quadprog`, e.g. the example passes B as the first argument to `quadprog`, so `qpbox4mult` must accept B as the first argument.

The second argument to the Hessian multiply function 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. Finally, any additional parameters passed to `quadprog` are passed to the Hessian multiply function, so `qpbox4mult` must accept the same additional parameters, e.g. the matrix A.

```
function W = qpbox4mult(B,Y,A);
%QPBOX4MULT Hessian matrix product with dense structured Hessian.
% W = qpbox4mult(B,Y,A) 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.
% 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);
```

Note `qpbox4mult` is a subfunction of `runqpbox4.m` in the Optimization Toolbox.

Step 3: Call a quadratic minimization routine with a starting point.

Load the problem parameters from `qpbox4.mat`. Use `optimset` to set the `HessMult` option to a function handle that points to `qpbox4mult`. Call `quadprog` with `B` as the first argument and `A` as an additional parameter.

```
load qpbox4          % Get xstart, u, l, B, A, f  
mtxmpy = @qpbox4mult; % Function handle to function qpbox4mult  
options = optimset('HessMult',mtxmpy);  
[x,fval,exitflag,output] = quadprog(B,f,[],[],[],[],l,u,...  
                                xstart,options,A);  
  
Optimization terminated successfully:  
Relative function value changing by less than  
sqrt(OPTIONS.TolFun), no negative curvature detected in Hessian  
this iteration, and the rate of progress (change in f(x)) is slow
```

After 18 iterations with a total of 30 PCG iterations, the function value is reduced to

```
fval =  
-1.0538e+003
```

and the firstorder optimality is

```
output.firstorderopt =  
0.0043
```

Note Type `[fval,exitflag,output] = runqpbox4` to run the code above and display the values for `fval`, `exitflag`, and `output`.

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.

Since the preconditioner is more approximate than when H is available explicitly, adjusting the `TolPcg` parameter to a somewhat smaller value may be required. This example is the same as the previous one, but reduces `TolPcg` from the default 0.1 to 0.01.

```
options = optimset('HessMult','mtxmpy','TolPcg',0.01);
[x,fval,exitflag,output]=quadprog(B,f,[],[],[],[],1,u,xstart,...
                                options,A);
```

```
Optimization terminated successfully:
Relative function value changing by less than
sqrt(OPTIONS.TolFun), no negative curvature detected in Hessian
this iteration, and the rate of progress (change in f(x)) is slow
```

After 18 iterations and 50 PCG iterations the function value has the same value to 5 significant digits

```
fval =
-1.0538e+003
```

but the first order optimality is further reduced.

```
output.firstorderopt =
0.0028
```

Note Decreasing `TolPcg` too much can substantially increase the number of PCG iterations.

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 =
    1
resnorm =
    22.5794
output =
    algorithm: 'large-scale: trust-region reflective Newton'
    firstorderopt: 2.7870e-005
    iterations: 10
    cgiterations: 42
```

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 3-8, and g is the gradient.

The first-order optimality can be improved (decreased) by using a sparse QR -factorization in each iteration: set `PrecondBandWidth` to `inf`.

```
options = optimset('PrecondBandWidth',inf);
[x,resnorm,residual,exitflag,output] = ...
    lsqlin(C,d,[],[],[],[],lb,[],[],options);
```

The number of iterations and the first-order optimality both decrease.

```
exitflag =
    1
resnorm =
    22.5794
```

```

output =
    algorithm: 'large-scale: trust-region reflective Newton'
    firstorderopt: 5.5907e-015
    iterations: 12
    cgiterations: 11

```

Linear Programming with Equalities and Inequalities

The problem is

$$\min f^T x \quad \text{such that} \quad \begin{aligned} A_{eq} \cdot x &= b_{eq} \\ A \cdot x &\leq b \\ x &\geq 0 \end{aligned}$$

and you can load the matrices and vectors A , A_{eq} , b , b_{eq} , 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.

You can use `linprog` to solve the problem.

```
[x,fval,exitflag,output] = ...
    linprog(f,A,b,Aeq,beq,lb,[],[],optimset('Display','iter'));
```

Since the iterative display was set using `optimset`, the results printed to the command line window are

Residuals:	Primal	Dual	Duality	Total
	Infeas	Infeas	Gap	Rel
	$A \cdot x - b$	$A' \cdot y + z - f$	$x' \cdot z$	Error

Iter 0:	1.50e+003	2.19e+001	1.91e+004	1.00e+002
Iter 1:	1.15e+002	2.94e-015	3.62e+003	9.90e-001
Iter 2:	1.16e-012	2.21e-015	4.32e+002	9.48e-001
Iter 3:	3.23e-012	5.16e-015	7.78e+001	6.88e-001
Iter 4:	5.78e-011	7.61e-016	2.38e+001	2.69e-001
Iter 5:	9.31e-011	1.84e-015	5.05e+000	6.89e-002
Iter 6:	2.96e-011	1.62e-016	1.64e-001	2.34e-003
Iter 7:	1.51e-011	2.74e-016	1.09e-005	1.55e-007
Iter 8:	1.51e-012	2.37e-016	1.09e-011	1.51e-013

Optimization terminated successfully.

For this problem the large-scale 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 =  
    1  
fval =  
-70.0000  
output =  
    iterations: 8  
    cgiterations: 0  
    algorithm: 'lipsol'
```

Linear Programming with Dense Columns in the Equalities

The problem is


$$\min f^T x \quad \text{such that} \quad \begin{array}{l} Aeq \cdot x = beq \\ lb \leq x \leq ub \end{array}$$

and you can load the matrices and vectors `Aeq`, `beq`, `f`, `lb`, and `ub` into the MATLAB workspace with

```
load densecolumns
```

The problem in `densecolumns.mat` has 1677 variables and 627 equalities with lower bounds on all the variables, and upper bounds on 399 of the variables. The equality matrix `Aeq` has dense columns among its first 25 columns, which is easy to see with a `spy` plot.

```
spy(Aeq)
```

You can use `linprog` to solve the problem.

```
[x,fval,exitflag,output] = ...
    linprog(f,[],[],Aeq,beq,lb,ub,[],optimset('Display','iter'));
```

Since the iterative display was set using `optimset`, the results printed to the command line window are

Residuals:	Primal Infeas A*x-b	Dual Infeas A'*y+z-w-f	Upper Bounds {x}+s-ub	Duality Gap x'*z+s'*w	Total Rel Error
Iter 0:	1.67e+003	8.11e+002	1.35e+003	5.30e+006	2.92e+001
Iter 1:	1.37e+002	1.33e+002	1.11e+002	1.27e+006	2.48e+000
Iter 2:	3.56e+001	2.38e+001	2.89e+001	3.42e+005	1.99e+000
Iter 3:	4.86e+000	8.88e+000	3.94e+000	1.40e+005	1.89e+000
Iter 4:	4.24e-001	5.89e-001	3.44e-001	1.91e+004	8.41e-001
Iter 5:	1.23e-001	2.02e-001	9.97e-002	8.41e+003	5.79e-001
Iter 6:	3.98e-002	7.91e-002	3.23e-002	4.05e+003	3.52e-001
Iter 7:	7.25e-003	3.83e-002	5.88e-003	1.85e+003	1.85e-001
Iter 8:	1.47e-003	1.34e-002	1.19e-003	8.12e+002	8.52e-002
Iter 9:	2.52e-004	3.39e-003	2.04e-004	2.78e+002	2.99e-002
Iter 10:	3.46e-005	1.08e-003	2.81e-005	1.09e+002	1.18e-002
Iter 11:	6.95e-007	1.53e-012	5.64e-007	1.48e+001	1.62e-003
Iter 12:	1.04e-006	2.26e-012	3.18e-008	8.32e-001	9.09e-005
Iter 13:	3.08e-006	1.23e-012	3.86e-009	7.26e-002	7.94e-006
Iter 14:	3.75e-007	1.09e-012	6.53e-012	1.11e-003	1.21e-007
Iter 15:	5.21e-008	1.30e-012	3.27e-013	8.62e-008	9.15e-010

Optimization terminated successfully.

You can see the returned values of `exitflag`, `fval`, and `output`.

```
exitflag =  
    1  
fval =  
    9.1464e+003  
output =  
    iterations: 15  
    cgiterations: 225  
    algorithm: 'lipsol'
```

This time the number of PCG iterations (in `output.cgiterations`) is nonzero because the dense columns in `Aeq` are detected. Instead of using a sparse Cholesky factorization, `linprog` tries to use the Sherman-Morrison formula to solve a linear system involving `Aeq*Aeq'`. If the Sherman-Morrison formula does not give a satisfactory residual, a PCG iteration is used. See the “Main Algorithm” section in “Large-Scale Linear Programming” on page 3-14.

Default Parameter Settings

The `options` structure contains parameters used in the optimization routines. If, on the first call to an optimization routine, the `options` structure is not provided, or is empty, a set of default parameters is generated. Some of the default options parameters are calculated using factors based on problem size, such as `MaxFunEvals`. Some parameters are dependent on the specific optimization routines and are documented on those function reference pages (See the “Function Reference” chapter).

Table 4-3, Optimization Options Parameters, on page 4-11 provides an overview of all the parameters in the `options` structure.

Changing the Default Settings

The function `optimset` creates or updates an `options` variable to pass to the various optimization functions. The arguments to the `optimset` function are parameter name and parameter value pairs, such as `TolX` and `1e-4`. Any unspecified properties have default values. You need to type only enough leading characters to define the parameter name uniquely. Case is ignored for parameter names. For parameter values that are strings, however, case and the exact string are necessary.

`help optimset` provides information that defines the different parameters and describes how to use them.

Here are some examples of the use of `optimset`.

Returning All Parameters

`optimset` returns all the parameters that can be set with typical values and default values.

Determining Parameters Used by a Function

The `options` structure defines the parameters that can be used by the functions provided by the toolbox. Because functions do not use all the parameters, it may be useful to find which parameters are used by a particular function.

To determine which `options` structure fields are used by a function, pass the name of the function (in this example, `fmincon`) to `optimset`

```
optimset('fmincon')
```

or

```
optimset fmincon
```

This statement returns a structure. Fields not used by the function have empty values (`[]`); fields used by the function are set to their default values for the given function.

Displaying Output

To display output at each iteration, enter

```
options = optimset('Display', 'iter');
```

This command sets the `options.Display` field value to `'iter'`, which causes the toolbox to display output at each iteration. You can also turn off any output display (`'off'`), display output only at termination (`'final'`), or display output only if the problem fails to converge (`'notify'`).

Running Medium-Scale Optimization

For functions that support medium- and large-scale optimization problems, the default is for the function to use the large-scale algorithm. To use the medium-scale algorithm, enter

```
options = optimset('LargeScale', 'off');
```

Setting More Than One Parameter

You can specify multiple parameters with one call to `optimset`. For example, to reset the output option and the tolerance on x , enter

```
options = optimset('Display', 'iter', 'TolX', 1e-6);
```

Updating an options Structure

To update an existing options structure, call `optimset` and pass options as the first argument.

```
options = optimset(options, 'Display', 'iter', 'TolX', 1e-6);
```

Retrieving Parameter Values

Use the `optimget` function to get parameter values from an options structure. For example, to get the current display option, enter

```
verbosity = optimget(options, 'Display');
```

Displaying Iterative Output

This section describes the column headings used in the iterative output of:

- Medium-scale algorithms
- Large-scale algorithms

Output Headings: Medium-Scale Algorithms

When the options `Display` parameter is set to `'iter'` for `fminsearch`, `fminbnd`, `fzero`, `fgoalattain`, `fmincon`, `lsqcurvefit`, `fminunc`, `fsolve`, `lsqnonlin`, `fminimax`, and `fseminf`, output is produced in column format.

`fminsearch`

For `fminsearch` the column headings are

```
Iteration   Func-count   min f(x)      Procedure
```

where:

- `Iteration` is the iteration number.
- `Func-count` is the number of function evaluations.
- `min f(x)` is the minimum function value in the current simplex.
- `Procedure` gives the current simplex operation: `initial`, `expand`, `reflect`, `shrink`, `contract inside` and `contract outside`.

`fzero` and `fminbnd`

For `fzero` and `fminbnd` the column headings are

```
Func-count   x           f(x)         Procedure
```

where:

- `Func-count` is the number of function evaluations (which for `fzero` is the same as the number of iterations).
- `x` is the current point.
- `f(x)` is the current function value at `x`.
- `Procedure` gives the current operation. For `fzero` these include `initial` (initial point), `search` (search for a interval containing a zero), `bisection`

(bisection search), and interpolation. For `fminbnd`, the possible operations are initial, golden (golden section search), and parabolic (parabolic interpolation).

fminunc

For `fminunc`, the column headings are

Iteration	Func-count	f(x)	Step-size	Directional derivative
-----------	------------	------	-----------	---------------------------

where:

- Iteration is the iteration number.
- Func-count is the number of function evaluations.
- $f(x)$ is the current function value.
- Step-size is the step-size in the current search direction.
- Directional derivative is the gradient of the function along the search direction.

fsolve, lsqnonlin, and lsqcurvefit

For `fsolve`, `lsqnonlin`, and `lsqcurvefit` the headings are

Iteration	Func-count	Residual	Step-size	Directional derivative	Lambda
-----------	------------	----------	-----------	---------------------------	--------

where Iteration, Func-count, Step-size, and Directional derivative are the same as for `fminunc`, and:

- Residual is the residual (sum-of-squares) of the function.
- Lambda is the λ_k value defined in “Least Squares Optimization” on page 2-19. (This value is printed when the Levenberg-Marquardt method is used and omitted when the Gauss-Newton method is used.)

fmincon and fsemif

For `fmincon` and `fsemif` the headings are

Iter	F-count	f(x)	max constraint	Step-size	Directional derivative	Procedure
------	---------	------	-------------------	-----------	---------------------------	-----------

where:

- Iter is the iteration number.
- F-count is the number of function evaluations.
- $f(x)$ is the current function value.
- max constraint is the maximum constraint violation.
- Step-size is the step-size in the search direction.
- Directional derivative is the gradient of the function along the search direction.
- Procedure gives a messages about the Hessian update and QP subproblem.

The Procedure messages are discussed in “Updating the Hessian Matrix” on page 2-30.

For fgoalattain and fminimax, the headings are the same as for fmincon except $f(x)$ and max constraint are combined into $\text{Max}\{F, \text{constraints}\}$. $\text{Max}\{F, \text{constraints}\}$ gives the maximum goal violation or constraint violation for fgoalattain, and the maximum function value or constraint violation for fminimax.

Output Headings: Large-Scale Algorithms

fminunc

For fminunc, the column headings are

Iteration	f(x)	Norm of step	First-order optimality	CG-iterations
-----------	------	-----------------	---------------------------	---------------

where:

- Iteration is the iteration number.
- $f(x)$ is the current function value.
- Norm of step is the norm of the current step-size.
- First-order optimality is the infinity norm of the current gradient.
- CG-iterations is the number of iterations taken by PCG (see “Preconditioned Conjugate Gradients” on page 3-6) at the current (optimization) iteration.

lsqnonlin, lsqcurvefit, and fsolve

For `lsqnonlin`, `lsqcurvefit`, and `fsolve` the column headings are

Iteration	Func-count	f(x)	Norm of step	First-order optimality	CG-iterations
-----------	------------	------	--------------	------------------------	---------------

where:

- Iteration is the iteration number.
- Func-count is the number of function evaluations.
- f(x) is the sum-of-the-squares of the current function values.
- Norm of step is the norm of the current step-size.
- First-order optimality is a measure of first-order optimality. 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 3-8 and g is the gradient. For unconstrained problems, it is the infinity norm of the current gradient.
- CG-iterations is the number of iterations taken by PCG (see “Preconditioned Conjugate Gradients” on page 3-6) at the current (optimization) iteration.

fmincon

For `fmincon`, the column headings are

Iteration	f(x)	Norm of step	First-order optimality	CG-iterations
-----------	------	--------------	------------------------	---------------

where:

- Iteration is the iteration number.
- f(x) is the current function value.
- Norm of step is the norm of the current step-size.
- First-order optimality is a measure of first-order optimality. 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 3-8 and g is the gradient. For equality constrained problems, it is the infinity norm of the projected gradient. (The projected gradient is the gradient projected into the nullspace of A_{eq} .)

- CG-iterations is the number of iterations taken by PCG (see “Preconditioned Conjugate Gradients” on page 3-6) at the current (optimization) iteration.

linprog

For linprog the column headings are

Residuals:	Primal	Dual	Upper	Duality	Total
	Infeas	Infeas	Bounds	Gap	Rel
	$A^*x - b$	$A' *y + z - w - f$	$\{x\} + s - ub$	$x' *z + s' *w$	Error

where:

- Primal Infeas $A^*x - b$ is the norm of the residual $A^*x - b$.
- Dual Infeas $A' *y + z - w - f$ is the norm of the residual $A' *y + z - w - f$ (where w is all zero if there are no finite upper bounds).
- Upper Bounds $\{x\} + s - ub$ is the norm of the residual $spones(s) .*x + s - ub$ (which is defined to be zero if all variables are unbounded above). This column is not printed if no finite upper bounds exist.
- Duality Gap $x' *z + s' *w$ is the duality gap (see “Large-Scale Linear Programming” on page 3-14) between the primal objective and the dual objective. s and w only appear in this equation if there are finite upper bounds.
- Total Rel Error is the total relative error described at the end of the “Main Algorithm” subsection of “Large-Scale Linear Programming” on page 3-14.

Optimization of Inline Objects Instead of M-Files

The routines in the Optimization Toolbox also perform optimization on inline objects, avoiding the need to write M-files to define functions.

To represent a mathematical function at the command line, create an inline object from a string expression. For example, you can create an inline object of the humps function (use the command `type humps` to see the M-file function `humps.m`).

```
f = inline('1./((x-0.3).^2 + 0.01) + 1./((x-0.9).^2 + 0.04)-6');
```

You can then evaluate `f` at 2.0.

```
f(2.0)
ans =
    -4.8552
```

And you can pass `f` to an optimization routine to minimize it.

```
x = fminbnd(f, 3, 4)
```

You can also create functions of more than one argument with `inline` by specifying the names of the input arguments along with the string expression. For example, to use `lsqcurvefit`, you need a function that takes two input arguments, `x` and `xdata`

```
f= inline('sin(x).*xdata +(x.^2).*cos(xdata)','x','xdata')
x = pi; xdata = pi*[4;2;3];
f(x, xdata)
ans =
    9.8696e+000
    9.8696e+000
   -9.8696e+000
```

and then call `lsqcurvefit`.

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

Other examples that use this technique:

- A matrix equation

```
x = fsolve(inline('x*x*x-[1,2;3,4]'),ones(2,2))
```

- **A nonlinear least squares problem**

```
x = lsqnonlin(inline('x*x-[3 5;9 10]'),eye(2,2))
```

- **An example using fgoalattain where the function has additional arguments to pass to the optimization routine. For example, if the function to be minimized has additional arguments A, B, and C,**

```
fun = inline('sort(eig(A+B*x*C))','x','A','B','C');  
x = fgoalattain(fun,-ones(2,2),[-5,-3,-1],[5, 3, 1],...  
[ ],[ ],[ ],[ ],-4*ones(2),4*ones(2),[ ],[ ],A,B,C);
```

solves the problem detailed on the fgoalattain reference page.

Typical Problems and How to Deal with Them

Optimization problems can take many iterations to converge and can be sensitive to numerical problems such as truncation and round-off error in the calculation of finite-difference gradients. Most optimization problems benefit from good starting guesses. This improves the execution efficiency and can help locate the global minimum instead of a local minimum.

Advanced problems are best solved by an evolutionary approach whereby a problem with a smaller number of independent variables is solved first. Solutions from lower order problems can generally be used as starting points for higher order problems by using an appropriate mapping.

The use of simpler cost functions and less stringent termination criteria in the early stages of an optimization problem can also reduce computation time. Such an approach often produces superior results by avoiding local minima.

The Optimization Toolbox functions can be applied to a large variety of problems. Used with a little “conventional wisdom,” many of the limitations associated with optimization techniques can be overcome. Additionally, problems that are not typically in the standard form can be handled by using an appropriate transformation. Below is a list of typical problems and recommendations for dealing with them.

Table 1-1: Troubleshooting

Problem	Recommendation
The solution does not appear to be a global minimum.	There is no guarantee that you have a global minimum unless your problem is continuous and has only one minimum. Starting the optimization from a number of different starting points may help to locate the global minimum or verify that there is only one minimum. Use different methods, where possible, to verify results.

Table 1-1: Troubleshooting (Continued)

Problem	Recommendation
<p>fminunc produces warning messages and seems to exhibit slow convergence near the solution.</p>	<p>If you are not supplying analytically determined gradients and the termination criteria are stringent, fminunc often exhibits slow convergence near the solution due to truncation error in the gradient calculation. Relaxing the termination criteria produces faster, although less accurate, solutions. For the medium-scale algorithm, another option is adjusting the finite-difference perturbation levels, DiffMinChange and DiffMaxChange, which may increase the accuracy of gradient calculations.</p>
<p>Sometimes an optimization problem has values of x for which it is impossible to evaluate the objective function fun or the nonlinear constraints function nonlcon.</p>	<p>Place bounds on the independent variables or make a penalty function to give a large positive value to f and g when infeasibility is encountered. For gradient calculation the penalty function should be smooth and continuous.</p>
<p>The function that is being minimized has discontinuities.</p>	<p>The derivation of the underlying method is based upon functions with continuous first and second derivatives. Some success may be achieved for some classes of discontinuities when they do not occur near solution points. One option is to smooth the function. For example, the objective function might include a call to an interpolation function to do the smoothing.</p> <p>Or, for the medium-scale algorithms, the finite-difference parameters may be adjusted in order to jump over small discontinuities. The variables DiffMinChange and DiffMaxChange control the perturbation levels for x used in the calculation of finite-difference gradients. The perturbation, Δx, is always in the range $\text{DiffMinChange} < \Delta x < \text{DiffMaxChange}$.</p>

Table 1-1: Troubleshooting (Continued)

Problem	Recommendation
Warning messages are displayed.	<p>This sometimes occurs when termination criteria are overly stringent, or when the problem is particularly sensitive to changes in the independent variables. This usually indicates truncation or round-off errors in the finite-difference gradient calculation, or problems in the polynomial interpolation routines. These warnings can usually be ignored because the routines continue to make steps toward the solution point; however, they are often an indication that convergence will take longer than normal. Scaling can sometimes improve the sensitivity of a problem.</p>
The independent variables, x , only can take on discrete values, for example, integers.	<p>This type of problem occurs commonly when, for example, the variables are the coefficients of a filter that are realized using finite-precision arithmetic or when the independent variables represent materials that are manufactured only in standard amounts.</p> <p>Although the Optimization Toolbox functions are not explicitly set up to solve discrete problems, some discrete problems can be solved by first solving an equivalent continuous problem. Discrete variables can be progressively eliminated from the independent variables, which are free to vary.</p> <p>Eliminate a discrete variable by rounding it up or down to the nearest best discrete value. After eliminating a discrete variable, solve a reduced order problem for the remaining free variables. Having found the solution to the reduced order problem, eliminate another discrete variable and repeat the cycle until all the discrete variables have been eliminated.</p> <p><code>dfildemo</code> is a demonstration routine that shows how filters with fixed precision coefficients can be designed using this technique.</p>

Table 1-1: Troubleshooting (Continued)

Problem	Recommendation
The minimization routine appears to enter an infinite loop or returns a solution that does not satisfy the problem constraints.	Your objective (<code>fun</code>), constraint (<code>nonlcon</code> , <code>seminfcon</code>), or gradient (computed by <code>fun</code>) functions may be returning <code>Inf</code> , <code>NaN</code> , or complex values. The minimization routines expect only real numbers to be returned. Any other values may cause unexpected results. Insert some checking code into the user-supplied functions to verify that only real numbers are returned (use the function <code>isfinite</code>).
You do not get the convergence you expect from the <code>lsqnonlin</code> routine.	You may be forming the sum of squares explicitly and returning a scalar value. <code>lsqnonlin</code> expects a vector (or matrix) of function values that are squared and summed internally.

Converting Your Code to Version 2 Syntax

Most of the function names and calling sequences have changed in Version 2 to accommodate new functionality and to clarify the roles of the input and output variables.

As a result, if you want to use the new versions of these functions, you need to modify any code that currently uses the old function names and calling sequences.

This table lists the functions provided by the Optimization Toolbox and indicates the functions whose names changed in Version 2.

Old (Version 1.5) Name	New (Version 2) Name
attgoal	fgoalattain
conls	lsqlin
constr	fmincon
curvefit	lsqcurvefit
fmin	fminbnd
fmins	fminsearch
fminu	fminunc
fsolve	fsolve (name unchanged)
fzero	fzero (name unchanged)
leastsq	lsqnonlin
minimax	fminimax
nls	lsqnonneg
lp	linprog
qp	quadprog
seminf	fseminf

This section:

- Tells you how to override default parameter settings with the new `optimset` and `optimget` functions
- Explains the reasons for the new calling sequences and explains how to convert your code.
- Provides a detailed example of rewriting a call to the `constr` function to call the new `fmincon` function instead.

In addition to the information in this section, consult the M-file help for the new functions for more information about the arguments they take. For example, to see the help for `fmincon`, type

```
help fmincon
```

Using `optimset` and `optimget`

The `optimset` function replaces `foptions` for overriding default parameter settings. See “Changing the Default Settings” on page 1-63 for more information on using `optimset` and `optimget`.

New Calling Sequences

Version 2 of the toolbox makes these changes in the calling sequences:

- Equality constraints and inequality constraints are now supplied as separate input arguments.
- Linear constraints and nonlinear constraints are now supplied as separate input arguments.
- The gradient of the objective is computed in the same function as the objective, rather than in a separate function, in order to provide more efficient computation (because the gradient and objective often share similar computations). Similarly, the gradient of the nonlinear constraints is computed by the (now separate) nonlinear constraint function.
- The Hessian matrix can be provided by the objective function. (This matrix is used only by the new large-scale algorithms.)
- Flags are *required* to indicate when extra information is available:
 - `options.GradObj = 'on'` indicates the user-supplied gradient of the objective function is available.

- `options.GradConstr = 'on'` indicates the user-supplied gradient of the constraints is available.
- `options.Hessian = 'on'` indicates the user-supplied Hessian of the objective function is available.
- Each function takes an options structure to adjust parameters to the optimization functions (see `optimset`, `optimget`).
- The new default output gives information upon termination (the old default was no output, the new default is `options.display = 'final'`).
- Each function returns an `exitflag` that denotes the termination state.
- The default uses the new large-scale methods when possible. If you want to use the older algorithms (referred to as medium-scale algorithms in other parts of this *User's Guide*), set `options.LargeScale = 'off'`.

Algorithm terminating conditions have been fine tuned. The stopping conditions relating to `TolX` and `TolFun` for the large-scale *and* medium-scale code are joined using OR instead of AND for these functions: `fgoalattain`, `fmincon`, `fminimax`, `fminunc`, `fseminf`, `fsolve`, and `lsqnonlin`. As a result, you may need to specify stricter tolerances; the defaults reflect this change.

Each function now has an output structure that contains information about the problem solution relevant to that function.

The `lambda` is now a structure where each field is the Lagrange multipliers for a type of constraint. For more information, see the individual functions in the “Function Reference” chapter.

The sections below describe how to convert from the old function names and calling sequences to the new ones. The calls shown are the most general cases, involving all possible input and output arguments. Note that many of these arguments are optional; see the online help for these functions for more information.

Converting from `attgoal` to `fgoalattain`

In Version 1.5, you used this call to `attgoal`

```
OPTIONS = foptions;  
[X,OPTIONS] = attgoal('FUN',x0,GOAL, WEIGHT, OPTIONS, VLB, VUB,  
    'GRADFUN', P1, P2,...);
```

with `[F] = FUN(X,P1,...)` and `[DF] = GRADFUN(X,P1,...)`.

In Version 2, you call `fgoalattain` like this

```

OPTIONS = optimset('fgoalattain');
[X,FVAL,ATTAINFACTOR,EXITFLAG,OUTPUT,LAMBDA] =
    fgoalattain(@FUN,x0,GOAL,WEIGHT,A,B,Aeq,Beq,VLB,VUB,
    @NONLCON,OPTIONS,P1,P2,...);
    
```

with `[F,DF] = FUN(X,P1,P2,...)` and `NONLCON = []`.

The `fgoalattain` function now allows nonlinear constraints, so you can now define

```
[Cineq,Ceq,DCineq,DCeq] = NONLCON(X,P1,...)
```

Converting from `conls` to `lsqlin`

In Version 1.5, you used this call to `conls`.

```
[X,LAMBDA,HOW] = conls(A,b,C,d,VLB,VUB,X0,N,DISPLAY);
```

In Version 2, convert the input arguments to the correct form for `lsqlin` by separating the equality and inequality constraints.

```

Ceq = C(1:N,:);
deq = d(1:N);
C = C(N+1:end,:);
d = d(N+1:end,:);
    
```

Now call `lsqlin` like this.

```

OPTIONS = optimset('Display','final');
[X,RESNORM,RESIDUAL,EXITFLAG,OUTPUT,LAMBDA] =
    lsqlin(A,b,C,d,Ceq,deq,VLB,VUB,X0,OPTIONS);
    
```

Converting from `constr` to `fmincon`

In Version 1.5, you used this call to `constr`

```

[X,OPTIONS,LAMBDA,HESS] =
    constr('FUN',x0,OPTIONS,VLB,VUB,'GRADFUN',P1,P2,...);
    
```

with `[F,C] = FUN(X,P1,...)` and `[G,DC] = GRADFUN(X,P1,...)`.

In Version 2, replace `FUN` and `GRADFUN` with two new functions:

- OBJFUN, which returns the objective function, the gradient (first derivative) of this function, and its Hessian matrix (second derivative).

```
[F,G,H] = OBJFUN(X,P1,...)
```

- NONLCON, which returns the functions for the nonlinear constraints (both inequality and equality constraints) and their gradients.

```
[C,Ceq,DC,DCEq] = NONLCON(X,P1,...)
```

Now call `fmincon` like this.

```
% OBJFUN supplies the objective gradient and Hessian;  
% NONLCON supplies the constraint gradient.  
OPTIONS =  
    optimset('GradObj','on','GradConstr','on','Hessian','on');  
[X,FVAL,EXITFLAG,OUTPUT,LAMBDA,GRAD,HESSIAN] =  
    fmincon(@OBJFUN,x0,A,B,Aeq,Beq,VLB,VUB,@NONLCON,OPTIONS,  
    P1,P2,...);
```

See “Example of Converting from `constr` to `fmincon`” on page 1-85 for a detailed example.

Converting from `curvefit` to `lsqcurvefit`

In Version 1.5, you used this call to `curvefit`

```
[X,OPTIONS,FVAL,JACOBIAN] =  
    curvefit('FUN',x0,XDATA,YDATA,OPTIONS,'GRADFUN',P1,P2,...);
```

with `F = FUN(X,P1,...)` and `G = GRADFUN(X,P1,...)`.

In Version 2, replace `FUN` and `GRADFUN` with a single function that returns both `F` and `J`, the objective function and the Jacobian. (The Jacobian is the transpose of the gradient.)

```
[F,J] = OBJFUN(X,P1,...)
```

Now call `lsqcurvefit` like this.

```
OPTIONS = optimset('Jacobian','on');           % Jacobian is supplied  
VLB = []; VUB = [];                           % New arguments not in curvefit  
[X,RESNORM,F,EXITFLAG,OUTPUT,LAMBDA,JACOB] =  
    lsqcurvefit(@OBJFUN,x0,XDATA,YDATA,VLB,VUB,OPTIONS,  
    P1,P2,...);
```

Converting from `fmin` to `fminbnd`

In Version 1.5, you used this call to `fmin`.

```
[X,OPTIONS] = fmin('FUN',x1,x2,OPTIONS,P1,P2,...);
```

In Version 2, you call `fminbnd` like this.

```
[X,FVAL,EXITFLAG,OUTPUT] = fminbnd(@FUN,x1,x2,...
                                   OPTIONS,P1,P2,...);
```

Converting from `fmins` to `fminsearch`

In Version 1.5, you used this call to `fmins`.

```
[X,OPTIONS] = fmins('FUN',x0,OPTIONS,[],P1,P2,...);
```

In Version 2, you call `fminsearch` like this.

```
[X,FVAL,EXITFLAG,OUTPUT] = fminsearch(@FUN,x0,...
                                       OPTIONS,P1,P2,...);
```

Converting from `fminu` to `fminunc`

In Version 1.5, you used this call to `fminu`.

```
[X,OPTIONS] = fminu('FUN',x0,OPTIONS,'GRADFUN',P1,P2,...);
```

with `F = FUN(X,P1, ...)` and `G = GRADFUN(X,P1, ...)`.

In Version 2, replace `FUN` and `GRADFUN` with a single function that returns both `F` and `G` (the objective function and the gradient).

```
[F,G] = OBJFUN(X,P1, ...)
```

(This function can also return the Hessian matrix as a third output argument.)

Now call `fminunc` like this.

```
OPTIONS = optimset('GradObj','on');           % Gradient is supplied
[X,FVAL,EXITFLAG,OUTPUT,GRAD,HESSIAN] =
    fminunc(@OBJFUN,x0,OPTIONS,P1,P2,...);
```

If you have an existing `FUN` and `GRADFUN` that you do not want to rewrite, you can pass them both to `fminunc` by placing them in a cell array.

```
OPTIONS = optimset('GradObj','on');           % Gradient is supplied
```

```
[X,FVAL,EXITFLAG,OUTPUT,GRAD,HESSIAN] =  
    fminunc(@FUN,@GRADFUN,x0,OPTIONS,P1,P2,...);
```

Converting to the new form of fsolve

In Version 1.5, you used this call to `fsolve`

```
[X,OPTIONS] =  
    fsolve('FUN',x0,XDATA,YDATA,OPTIONS,'GRADFUN',P1,P2,...);
```

with `F = FUN(X,P1,...)` and `G = GRADFUN(X,P1,...)`.

In Version 2, replace `FUN` and `GRADFUN` with a single function that returns both `F` and `G` (the objective function and the gradient).

```
[F,G] = OBJFUN(X,P1, ...)
```

Now call `fsolve` like this.

```
OPTIONS = optimset('GradObj','on');           % Gradient is supplied  
[X,FVAL,EXITFLAG,OUTPUT,JACOBIAN] =  
    fsolve(@OBJFUN,x0,OPTIONS,P1,P2,...);
```

If you have an existing `FUN` and `GRADFUN` that you do not want to rewrite, you can pass them both to `fsolve` by placing them in a cell array.

```
OPTIONS = optimset('GradObj','on');           % Gradient is supplied  
[X,FVAL,EXITFLAG,OUTPUT,JACOBIAN] =  
    fsolve({@FUN,@GRADFUN},x0,OPTIONS,P1,P2,...);
```

Converting to the new form of fzero

In Version 1.5, you used this call to `fzero`.

```
X = fzero('F',X,TOL,TRACE,P1,P2,...);
```

In Version 2, replace the `TRACE` and `TOL` arguments with

```
if TRACE == 0,  
    val = 'none';  
elseif TRACE == 1  
    val = 'iter';  
end  
OPTIONS = optimset('Display',val,'TolX',TOL);
```

Now call `fzero` like this.

```
[X,FVAL,EXITFLAG,OUTPUT] = fzero(@F,X,OPTIONS,P1,P2,...);
```

Converting from `leastsq` to `lsqnonlin`

In Version 1.5, you used this call to `leastsq`

```
[X,OPTIONS,FVAL,JACOBIAN] =
    leastsq('FUN',x0,OPTIONS,'GRADFUN',P1,P2,...);
```

with `F = FUN(X,P1,...)` and `G = GRADFUN(X,P1, ...)`.

In Version 2, replace `FUN` and `GRADFUN` with a single function that returns both `F` and `J`, the objective function and the Jacobian. (The Jacobian is the transpose of the gradient.)

```
[F,J] = OBJFUN(X,P1, ...)
```

Now call `lsqnonlin` like this.

```
OPTIONS = optimset('Jacobian','on');          % Jacobian is supplied
VLB = []; VUB = [];                          % New arguments not in leastsq
[X,RESNORM,F,EXITFLAG,OUTPUT,LAMBDA,JACOBIAN] =
    lsqnonlin(@OBJFUN,x0,VLB,VUB,OPTIONS,P1,P2,...);
```

Converting from `lp` to `linprog`

In Version 1.5, you used this call to `lp`.

```
[X,LAMBDA,HOW] = lp(f,A,b,VLB,VUB,X0,N,DISPLAY);
```

In Version 2, convert the input arguments to the correct form for `linprog` by separating the equality and inequality constraints.

```
Aeq = A(1:N,:);
beq = b(1:N);
A = A(N+1:end,:);
b = b(N+1:end,:);
if DISPLAY
    val = 'final';
else
    val = 'none';
end
OPTIONS = optimset('Display',val);
```

Now call `linprog` like this.

```
[X,FVAL,EXITFLAG,OUTPUT,LAMBDA] =  
    linprog(f,A,b,Aeq,beq,VLB,VUB,X0,OPTIONS);
```

Converting from `minimax` to `fminimax`

In Version 1.5, you used this call to `minimax`

```
[X,OPTIONS] =  
    minimax('FUN',x0,OPTIONS,VLB,VUB,'GRADFUN',P1,P2,...);
```

with `F = FUN(X,P1,...)` and `G = GRADFUN(X,P1,...)`.

In Version 2, you call `fminimax` like this

```
OPTIONS = optimset('fminimax');  
[X,FVAL,MAXFVAL,EXITFLAG,OUTPUT,LAMBDA] =  
    fminimax(@OBJFUN,x0,A,B,Aeq,Beq,VLB,VUB,@NONLCON,OPTIONS,  
    P1,P2,...);
```

with `[F,DF] = OBJFUN(X,P1,...)` and
`[Cineq,Ceq,DCineq,DCEq] = NONLCON(X,P1,...)`.

Converting from `nls` to `lsqnonneg`

In Version 1.5, you used this call to `nls`.

```
[X,LAMBDA] = nls(A,b,tol);
```

In Version 2, replace the `tol` argument with

```
OPTIONS = optimset('Display','none','TolX',tol);
```

Now call `lsqnonneg` like this.

```
[X,RESNORM,RESIDUAL,EXITFLAG,OUTPUT,LAMBDA] =  
    lsqnonneg(A,b,X0,OPTIONS);
```

Converting from `qp` to `quadprog`

In Version 1.5, you used this call to `qp`.

```
[X,LAMBDA,HOW] = qp(H,f,A,b,VLB,VUB,X0,N,DISPLAY);
```

In Version 2, convert the input arguments to the correct form for `quadprog` by separating the equality and inequality constraints.


```

Aeq = A(1:N,:);
beq = b(1:N);
A = A(N+1:end,:);
b = b(N+1:end,:);
if DISPLAY
    val = 'final';
else
    val = 'none';
end
OPTIONS = optimset('Display',val);
    
```

Now call `quadprog` like this.

```

[X,FVAL,EXITFLAG,OUTPUT,LAMBDA] =
    quadprog(H,f,A,b,Aeq,beq,VLB,VUB,X0,OPTIONS);
    
```

Converting from `seminf` to `fseminf`

In Version 1.5, you used this call to `seminf`

```

[X,OPTIONS] = seminf('FUN',N,x0,OPTIONS,VLB,VUB,P1,P2,...);
    
```

with `[F,C,PHI1,PHI2,...,PHIN,S] = FUN(X,S,P1,P2,...)`.

In Version 2, call `fseminf` like this.

```

[X,FVAL,EXITFLAG,OUTPUT,LAMBDA] =
    fseminf(@OBJFUN,x0,N,@NONLCON,A,B,Aeq,Beq,VLB,VUB,OPTIONS,
    P1,P2,...);
    
```

with `F = OBJFUN(X,P1,...)` and

`[Cineq,Ceq,PHI1,PHI2,...,PHIN,S] = NONLCON(X,S,P1,...)`.

Example of Converting from `constr` to `fmincon`

Old Call to `constr`

```

OPTIONS = foptions;
OPTIONS(13) = 2;           % Two equality constraints
OPTIONS(1) = 1;
OPTIONS(9) = 1;
A1 = [ 1 4 -3]; b1 = 2;
A2 = [ 2 5 0]; b2 = 9;
    
```

```

x0 = [1; .5; .8];
LB = []; UB = [];
[X,OPTIONS,LAMBDA,HESS] = ...
    constr('myfuncon',x0,OPTIONS,LB,UB,'mygradcon',A1,b1,A2,b2);

```

```

% myfuncon.m
[F, C] = myfuncon(x,A1,b1,A2,b2)
F = x(1) + 0.0009*x(2)^3 + sin(x(3));

C(1,1) = A1*x-b;           % Equality linear constraint
C(2,1) = 3*x(1)^2-1;      % Equality nonlinear constraint

C(3,1) = A2*x-b2;         % Inequality linear constraint
C(4,1) = 7*sin(x(2))-1;   % Inequality nonlinear constraint

```

```

% mygradcon.m
[G, DC] = mygradcon(x,alpha)
G = [1;                    % Gradient of the objective
     3*0.0009*x(2)^2;
     cos(x(3))];

```

```

DC(:,1) = A1';           % Gradient of the constraints
DC(:,2) = [6*x(1); 0; 0];
DC(:,3) = A2';
DC(:,4) = [0; 7*cos(x(2)); 0];

```

New Call to fmincon

```

OPTIONS = optimset(...
    'Display', 'iter', ...
    'GradCheck', 'on', ... % Check gradients.
    'GradObj', 'on', ... % Gradient of objective is provided.
    'GradConstr', 'on'); % Gradient of constraints is provided.

```

```

A1 = [ 1 4 -3]; b1 = 2; % Linear equalities
A2 = [ 2 5 0]; b2 = 9; % Linear inequalities

```

```

x0 = [1; .5; .8];
LB = []; UB = [];

```

```
[X,FVAL,EXITFLAG,OUTPUT,LAMBDA,GRAD,HESSIAN] = ...
    fmincon(@myfun,x0,A2,b2,A1,b1,LB,UB,@mycon,OPTIONS);

% myfun.m
function [F,G] = myfun(x)
F = x(1) + 0.0009*x(2)^3 + sin(x(3));
G = [1;
     3*0.0009*x(2)^2;
     cos(x(3))];

% mycon.m
function [C,Ceq,DC,DCEq]= mycon(x)
Ceq(1,1) = 3*x(1)^2-1;           % Equality nonlinear constraint
C(1,1) = 7*sin(x(2))-1;        % Inequality nonlinear constraint
DCEq(:,1) = [6*x(1); 0; 0];    % Gradient of equality
                                % nonlinear constraint
DC(:,1) = [0; 7*cos(x(2)); 0]; % Gradient of inequality
                                % nonlinear constraint
```

Selected Bibliography

[1] Hairer, E., S. P. Norsett, and G. Wanner, *Solving Ordinary Differential Equations I – Nonstiff Problems*, Springer-Verlag, pages 183-184.

Standard Algorithms

Overview	2-2
Optimization Overview	2-4
Unconstrained Optimization	2-5
Quasi-Newton Methods	2-7
Line Search	2-9
Quasi-Newton Implementation	2-12
Hessian Update	2-12
Line Search Procedures	2-12
Least Squares Optimization	2-19
Gauss-Newton Method	2-20
Levenberg-Marquardt Method	2-22
Nonlinear Least Squares Implementation	2-24
Gauss-Newton Implementation	2-24
Levenberg-Marquardt Implementation	2-24
Constrained Optimization	2-26
Sequential Quadratic Programming (SQP)	2-27
Quadratic Programming (QP) Subproblem	2-28
SQP Implementation	2-30
Updating the Hessian Matrix	2-30
Quadratic Programming Solution	2-31
Line Search and Merit Function	2-35
Multiobjective Optimization	2-36
Introduction	2-36
Goal Attainment Method	2-42
Algorithm Improvements for Goal Attainment Method	2-43
Selected Bibliography	2-46

Overview

This chapter provides an introduction to the different optimization problem formulations and describes the *medium-scale* (i.e., standard) algorithms used in the toolbox functions. These algorithms have been chosen for their robustness and iterative efficiency. The choice of problem formulation (e.g., unconstrained, least squares, constrained, minimax, multiobjective, or goal attainment) depends on the problem being considered and the required execution efficiency.

Note Medium-scale is not a standard term and is used here only to differentiate these algorithms from the large-scale algorithms described in the “Large-Scale Algorithms” chapter.

This chapter is organized into these major topics.

Optimization Overview

Introduces optimization as a way of finding a set of parameters that can in some way be defined as optimal. These parameters are obtained by minimizing or maximizing an objective function, subject to equality or inequality constraints and/or parameter bounds.

Unconstrained Optimization

Discusses the use of quasi-Newton and line search methods for unconstrained optimization.

Quasi-Newton Implementation

Provides implementation details for the Hessian update and line search phases of the quasi-Newton algorithm used in `fminunc`.

Least Squares Optimization

Discusses the use of the Gauss-Newton and Levenberg-Marquardt methods for nonlinear least squares (LS) optimization.

Nonlinear Least Squares Implementation

Provides implementation details for the Gauss-Newton and Levenberg-Marquardt methods used in the nonlinear least squares optimization routines, `lsqnonlin` and `lsqcurvefit`.

Constrained Optimization

Discusses the use of the Kuhn-Tucker (KT) equations as the basis for Sequential Quadratic Programming (SQP) methods.

SQP Implementation

Provides implementation details for the Hessian matrix update, quadratic programming problem solution, and line search and merit function calculation phases of the SQP algorithm used in `fmincon`, `fminimax`, `fgoalattain`, and `fseminf`.

Multiobjective Optimization

Introduces multiobjective optimization and discusses strategies for dealing with competing objectives. It discusses in detail the use of the Goal Attainment method, and suggests improvements to the SQP method for use with the Goal Attainment method.

Selected Bibliography

Lists published materials that support concepts implemented in the medium-scale algorithms.

Optimization 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 may 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, may 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

$$\begin{array}{ll} \text{minimize} & f(x) \\ & x \in \mathfrak{R}^n \end{array} \quad (2-1)$$

subject to

$$\begin{array}{ll} G_i(x) = 0, & i = 1, \dots, m_e \\ G_i(x) \leq 0, & i = m_e + 1, \dots, m \\ x_l \leq x \leq x_u \end{array}$$

where x is the vector of design parameters, ($x \in \mathfrak{R}^n$), $f(x)$ is the objective function that returns a scalar value ($f(x): \mathfrak{R}^n \rightarrow \mathfrak{R}$), and the vector function $G(x)$ returns the values of the equality and inequality constraints evaluated at x ($G(x): \mathfrak{R}^n \rightarrow \mathfrak{R}^m$).

An efficient and accurate solution to this problem is not only dependent 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 may 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 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 very nonlinear 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 since 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

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

The minimum of this function is at $x = [1,1]$ where $f(x) = 0$. A contour map of this function is shown in Figure 2-1, 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 zig-zagging from one side of the valley to another. Note that towards the center of the plot, a number of larger steps are taken when a point lands exactly at the center of the valley.




Figure 2-1: Steepest Descent Method on Rosenbrock's Function (Eq. 2-2)

This type of function (Eq. 2-2), also known as the banana function, is notorious in unconstrained examples because of the way the curvature bends around the origin. Eq. 2-2 is used throughout this section to illustrate the use of a variety of optimization techniques. The contours have been plotted in exponential increments due to the steepness of the slope surrounding the U-shaped valley.

This section continues with discussions of the following methods for unconstrained optimization:

- Quasi-Newton Methods
- Line Search

"Quasi-Newton Implementation" on page 2-12 discusses implementation of these methods in the `fminunc` function.

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

$$\min_x \frac{1}{2} x^T H x + c^T x + b \quad (2-3)$$

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.,

$$\nabla f(x^*) = Hx^* + c = 0 \quad (2-4)$$

The optimal solution point, x^* , can be written as

$$x^* = -H^{-1}c \quad (2-5)$$

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. Generally, the formula of Broyden [3], Fletcher [12], Goldfarb [20], and Shanno [35] (BFGS) is thought to be the most effective for use in a general purpose method.

The formula given by BFGS is

$$H_{k+1} = H_k + \frac{q_k q_k^T}{q_k^T s_k} - \frac{H_k^T s_k s_k^T H_k}{s_k^T H_k s_k} \quad (2-6)$$

where

$$s_k = x_{k+1} - x_k$$

$$q_k = \nabla f(x_{k+1}) - \nabla f(x_k)$$

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 in which the direct inversion of H is avoided 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 [8], Fletcher, and Powell [14]. This uses the same formula as the above BFGS method (Eq. 2-6) 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

$$d = -H_k^{-1} \cdot \nabla f(x_k) \quad (2-7)$$

The quasi-Newton method is illustrated by the solution path on Rosenbrock's function (Eq. 2-2) in Figure 2-2, 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 2-2: BFGS Method on Rosenbrock's Function

Line Search

Most unconstrained and constrained methods use the solution of a subproblem to yield a search direction in which the solution is estimated to lie. The minimum along the line formed from this search direction is generally approximated using a search procedure (e.g., Fibonacci, Golden Section) or by a polynomial method involving interpolation or extrapolation (e.g., quadratic, cubic). Polynomial methods approximate a number of points with a univariate polynomial whose minimum can be calculated easily. Interpolation refers to the condition that the minimum is bracketed (i.e., the minimum lies in the area spanned by the available points), whereas extrapolation refers to a minimum located outside the range spanned by the available points. Extrapolation methods are generally considered unreliable for estimating minima for nonlinear functions. However, they are useful for estimating step length when

trying to bracket the minimum as shown in “Line Search Procedures” on page 2-12. Polynomial interpolation methods are generally the most effective in terms of efficiency when the function to be minimized is continuous. The problem is to find a new iterate x_{k+1} of the form

$$x_{k+1} = x_k + \alpha^* d \quad (2-8)$$

where x_k denotes the current iterate, d the search direction obtained by an appropriate method, and α^* is a scalar step length parameter that is the distance to the minimum.

Quadratic Interpolation

Quadratic interpolation involves a data fit to a univariate function of the form

$$m_q(\alpha) = a\alpha^2 + b\alpha + c \quad (2-9)$$

where an extremum occurs at a step length of

$$\alpha^* = \frac{-b}{2a} \quad (2-10)$$

This point may be a minimum or a maximum. It is a minimum when interpolation is performed (i.e., using a bracketed minimum) or when a is positive. Determination of coefficients, a and b , can be found using any combination of three gradient or function evaluations. It may also be carried out with just two gradient evaluations. The coefficients are determined through the formulation and solution of a linear set of simultaneous equations. Various simplifications in the solution of these equations can be achieved when particular characteristics of the points are used. For example, the first point can generally be taken as $\alpha = 0$. Other simplifications can be achieved when the points are evenly spaced. A general problem formula is as follows.

Given three unevenly spaced points $\{x_1, x_2, x_3\}$ and their associated function values $\{f(x_1), f(x_2), f(x_3)\}$ the minimum resulting from a second-order fit is given by

$$x_{k+1} = \frac{1}{2} \frac{\beta_{23}f(x_1) + \beta_{31}f(x_2) + \beta_{12}f(x_3)}{\gamma_{23}f(x_1) + \gamma_{31}f(x_2) + \gamma_{12}f(x_3)} \quad (2-11)$$

where

$$\beta_{ij} = x_i^2 - x_j^2$$

$$\gamma_{ij} = x_i - x_j$$

For interpolation to be performed, as opposed to extrapolation, the minimum must be bracketed so that the points can be arranged to give

$$f(x_2) < f(x_1) \text{ and } f(x_2) < f(x_3)$$

Cubic Interpolation

Cubic interpolation is useful when gradient information is readily available or when more than three function evaluations have been calculated. It involves a data fit to the univariate function

$$m_c(\alpha) = a\alpha^3 + b\alpha^2 + c\alpha + d \quad (2-12)$$

where the local extrema are roots of the quadratic equation

$$3a\alpha^2 + 2b\alpha + c = 0$$

To find the minimum extremum, take the root that gives $6a\alpha + 2b$ as positive. Coefficients a and b can be determined using any combination of four gradient or function evaluations, or alternatively, with just three gradient evaluations.

The coefficients are calculated by the formulation and solution of a linear set of simultaneous equations. A general formula, given two points, $\{x_1, x_2\}$, their corresponding gradients with respect to x , $\{\nabla f(x_1), \nabla f(x_2)\}$, and associated function values, $\{f(x_1), f(x_2)\}$ is

$$x_{k+1} = x_2 - (x_2 - x_1) \frac{\nabla f(x_2) + \beta_2 - \beta_1}{\nabla f(x_2) - \nabla f(x_1) + 2\beta_2} \quad (2-13)$$

where

$$\beta_1 = \nabla f(x_1) + \nabla f(x_2) - 3 \frac{f(x_1) - f(x_2)}{x_1 - x_2}$$

$$\beta_2 = (\beta_1^2 - \nabla f(x_1) \nabla f(x_2))^{1/2}.$$

Quasi-Newton Implementation

A quasi-Newton algorithm is used in `fminunc`. The algorithm consists of two phases:

- Determination of a direction of search (Hessian update)
- Line search procedures

Implementation details of the two phases are discussed below.

Hessian Update

The direction of search is determined by a choice of either the BFGS (Eq. 2-6) or the DFP method given in “Quasi-Newton Methods” on page 2-7 (set the options parameter `HessUpdate` to 'dfp' 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. Positive definiteness of H is achieved by ensuring that H is initialized to be positive definite and thereafter $q_k^T s_k$ (from Eq. 2-14) 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,

$$q_k^T s_k = \alpha_k (\nabla f(x_{k+1})^T d - \nabla f(x_k)^T d) \quad (2-14)$$

The condition that $q_k^T s_k$ is positive is always achieved by ensuring that a sufficiently accurate line search is performed. This is because the search direction, d , is a descent direction so that α_k and $-\nabla f(x_k)^T d$ 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.

Line Search Procedures

Two line search strategies are used depending on whether gradient information is readily available or whether it must be calculated using a finite difference method. When gradient information is available, the default is to use a cubic polynomial method. When gradient information is not available, the default is to use a mixed quadratic and cubic polynomial method.

Cubic Polynomial Method

In the proposed cubic polynomial method, a gradient and a function evaluation is made at every iteration, k . At each iteration an update is performed when a new point is found, x_{k+1} , which satisfies the condition that

$$f(x_{k+1}) < f(x_k) \quad (2-15)$$

At each iteration a step, α_k , is attempted to form a new iterate of the form

$$x_{k+1} = x_k + \alpha_k d \quad (2-16)$$

If this step does not satisfy the condition (Eq. 2-15) then α_k is reduced to form a new step, α_{k+1} . The usual method for this reduction is to use bisection (i.e., to continually halve the step length until a reduction is achieved in $f(x)$). However, this procedure is slow when compared to an approach that involves using gradient and function evaluations together with cubic interpolation/extrapolation methods to identify estimates of step length.

When a point is found that satisfies the condition (Eq. 2-15), an update is performed if $q_k^T s_k$ is positive. If it is not, then further cubic interpolations are performed until the univariate gradient term $\nabla f(x_{k+1})^T d$ is sufficiently small so that $q_k^T s_k$ is positive.

It is usual practice to reset α_k to unity after every iteration. However, note that the quadratic model (Eq. 2-3) is generally only a good one near to the solution point. Therefore, α_k , is modified at each major iteration to compensate for the case when the approximation to the Hessian is monotonically increasing or decreasing. To ensure that, as x_k approaches the solution point, the procedure reverts to a value of α_k close to unity, the values of $q_k^T s_k - \nabla f(x_k)^T d$ and α_{k+1} are used to estimate the closeness to the solution point and thus to control the variation in α_k .



Cubic Polynomial Line Search Procedures

After each update procedure, a step length α_k is attempted, following which a number of scenarios are possible. Consideration of all the possible cases is quite complicated and so they are represented pictorially below.

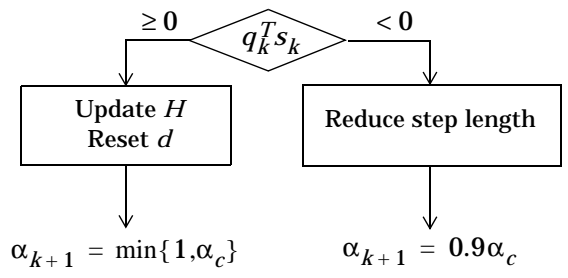

For each case:

- The left-hand point on the graph represents the point x_k .
- The slope of the line bisecting each point represents the slope of the univariate gradient, $\nabla f(x_k)^T d$, which is always negative for the left-hand point.
- The right-hand point is the point x_{k+1} after a step of α_k is taken in the direction d .

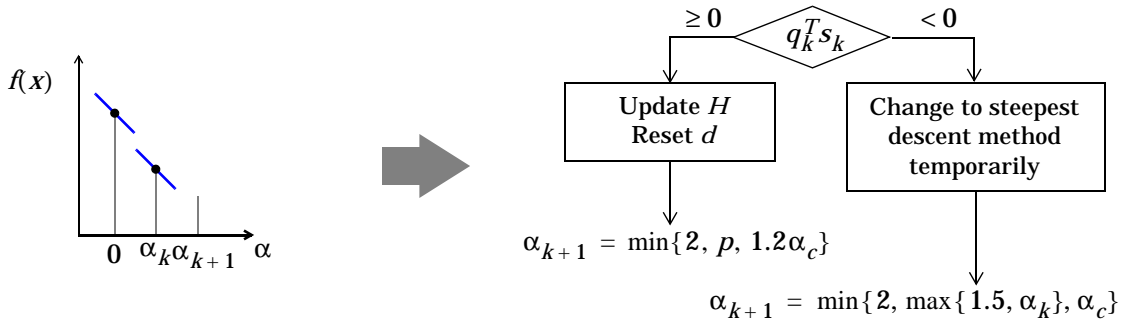
Case 1: $f(x_{k+1}) > f(x_k), \nabla f(x_{k+1})^T d > 0$




Case 2: $f(x_{k+1}) \leq f(x_k), \nabla f(x_{k+1})^T d \geq 0$



Case 3: $f(x_{k+1}) < f(x_k)$, $\nabla f(x_{k+1})^T d < 0$



Case 4: $f(x_{k+1}) \geq f(x_k)$, $\nabla f(x_{k+1})^T d \leq 0$ where $p = 1 + q_k^T s_k - \nabla f(x_{k+1})^T d + \min\{0, \alpha_{k+1}\}$



Cases 1 and 2 show the procedures performed when the value $\nabla f(x_{k+1})^T d$ is positive. Cases 3 and 4 show the procedures performed when the value $\nabla f(x_{k+1})^T d$ is negative. The notation $\min\{a, b, c\}$ refers to the smallest value of the set $\{a, b, c\}$.

At each iteration a cubically interpolated step length α_c is calculated and then used to adjust the step length parameter α_{k+1} . Occasionally, for very nonlinear functions α_c may be negative, in which case α_c is given a value of $2\alpha_k$. The methods for changing the step length have been refined over a period of time by considering a large number of test problems.

Certain robustness measures have also been included so that, even in the case when false gradient information is supplied, a reduction in $f(x)$ can be achieved

by taking a negative step. This is done by setting $\alpha_{k+1} = -\alpha_k/2$ when α_k falls below a certain threshold value (e.g., $1e-8$). This is important when extremely high precision is required if only finite difference gradients are available.

Mixed Cubic/Quadratic Polynomial Method

The cubic interpolation/extrapolation method has proved successful for a large number of optimization problems. However, when analytic derivatives are not available, the evaluating finite difference gradients is computationally expensive. Therefore, another interpolation/extrapolation method is implemented so that gradients are not needed at every iteration. The approach in these circumstances, when gradients are not readily available, is to use a quadratic interpolation method. The minimum is generally bracketed using some form of bisection method. This method, however, has the disadvantage that all the available information about the function is not used. For instance, a gradient calculation is always performed at each major iteration for the Hessian update. Therefore, given three points that bracket the minimum, it is possible to use cubic interpolation, which is likely to be more accurate than using quadratic interpolation. Further efficiencies are possible if, instead of using bisection to bracket the minimum, extrapolation methods similar to those used in the cubic polynomial method are used.


Hence, the method that is used in `fminunc`, `lsqnonlin`, `lsqcurvefit`, and `fsolve` is to find three points that bracket the minimum and to use cubic interpolation to estimate the minimum at each line search. The estimation of step length, at each minor iteration, j , is shown in the graphs below for a number of point combinations. The left-hand point in each graph represents the function value $f(x_1)$ and univariate gradient $\nabla f(x_k)$ obtained at the last update. The right-hand points represent the points accumulated in the minor iterations of the line search procedure.

The terms α_q and α_c refer to the minimum obtained from a respective quadratic and cubic interpolation or extrapolation. For highly nonlinear functions, α_c and α_q may be negative, in which case they are set to a value of $2\alpha_k$ so that they are always maintained to be positive. Cases 1 and 2 use quadratic interpolation with two points and one gradient to estimate a third point that brackets the minimum. If this fails, cases 3 and 4 represent the possibilities for changing the step length when at least three points are available.

When the minimum is finally bracketed, cubic interpolation is achieved using one gradient and three function evaluations. If the interpolated point is greater than any of the three used for the interpolation, then it is replaced with the point with the smallest function value. Following the line search procedure the Hessian update procedure is performed as for the cubic polynomial line search method.

The following graphs illustrate the line search procedures for cases 1 through 4, with a gradient only for the first point.

Case 1: $f(x_j) \geq f(x_k)$




Reduce step length



$$\alpha_{j+1} = \alpha_q$$

Case 2: $f(x_j) < f(x_k)$




Increase step length



$$\alpha_{j+1} = 1.2\alpha_q$$

Case 3: $f(x_{j+1}) < f(x_k)$




Increase step length



$$\alpha_{j+2} = \max\{1.2\alpha_q, 2\alpha_{j+1}\}$$

Case 4: $f(x_{j+1}) > f(x_k)$



Reduce step length

$\alpha_{j+2} = \alpha_c$

Least Squares Optimization

The line search procedures used in conjunction with a quasi-Newton method are used in the function `fminunc`. They are also used as part of a nonlinear least squares (LS) optimization routines, `lsqnonlin` and `lsqcurvefit`. In the least squares problem a function, $f(x)$ is minimized that is a sum of squares.

$$\min_{x \in \mathfrak{R}^n} f(x) = \frac{1}{2} \| F(x) \|_2^2 = \frac{1}{2} \sum_i F_i(x)^2 \quad (2-17)$$

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, $\phi(t)$, for vector x and scalar t . This problem can be expressed as

$$\min_{x \in \mathfrak{R}^n} \int_{t_2}^{t_1} (y(x, t) - \phi(t))^2 dt \quad (2-18)$$

where $y(x, t)$ and $\phi(t)$ are scalar functions.

When the integral is discretized using a suitable quadrature formula, Eq. 2-18 can be formulated as a least squares problem

$$\min_{x \in \mathfrak{R}^n} f(x) = \sum_{i=1}^m (\bar{y}(x, t_i) - \bar{\phi}(t_i))^2 \quad (2-19)$$

where \bar{y} and $\bar{\phi}$ include the weights of the quadrature scheme. Note that in this problem the vector $F(x)$ is

$$F(x) = \begin{bmatrix} \bar{y}(x, t_1) - \bar{\phi}(t_1) \\ \bar{y}(x, t_2) - \bar{\phi}(t_2) \\ \dots \\ \bar{y}(x, t_m) - \bar{\phi}(t_m) \end{bmatrix}$$

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 (Eq. 2-18) can be minimized using a

general unconstrained minimization technique as described in “Unconstrained Optimization” on page 2-5, 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 (Eq. 2-18) have a special structure.

Denoting the $m \times n$ 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 $H_i(x)$, we have

$$\begin{aligned} G(x) &= 2J(x)^T F(x) \\ H(x) &= 2J(x)^T J(x) + 2Q(x) \end{aligned} \tag{2-20}$$

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.

This section continues with discussions of the following methods for least squares optimization:

- Gauss-Newton Method
- Levenberg-Marquardt Method

“Nonlinear Least Squares Implementation” on page 2-24 discusses implementation of these methods in the `lsqnonlin` and `lsqcurvefit` functions.

Gauss-Newton Method

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.

$$\min_{x \in \mathfrak{R}^n} \| J(x_k) d_k - F(x_k) \|_2^2 \tag{2-21}$$

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.

To consider the efficiencies that are possible with the Gauss-Newton method. Figure 2-3, Gauss-Newton Method on Rosenbrock's Function, shows the path to the minimum on Rosenbrock's function (Eq. 2-2) when posed as a least squares problem. The Gauss-Newton method converges after only 48 function evaluations using finite difference gradients compared to 140 iterations using an unconstrained BFGS method.

The Gauss-Newton method often encounters problems when the second order term $Q(x)$ in Eq. 2-20 is significant. A method that overcomes this problem is the Levenberg-Marquardt method.




Figure 2-3: Gauss-Newton Method on Rosenbrock's Function

Levenberg-Marquardt Method

The Levenberg-Marquardt [25],[27] method uses a search direction that is a solution of the linear set of equations

$$(J(x_k)^T J(x_k) + \lambda_k I) d_k = -J(x_k) F(x_k) \quad (2-22)$$

where the scalar λ_k controls both the magnitude and direction of d_k . 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 a vector of zeros and a steepest descent direction. 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. This is illustrated in Figure 2-4, Levenberg-Marquardt Method on Rosenbrock's Function below. The solution for Rosenbrock's function (Eq. 2-2) 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 occasional poorer efficiency of the Levenberg-Marquardt method is compensated for by its increased robustness.




Figure 2-4: Levenberg-Marquardt Method on Rosenbrock's Function

Nonlinear Least Squares Implementation

For a general survey of nonlinear least squares methods, see Dennis [9]. Specific details on the Levenberg-Marquardt method can be found in Moré [28]. Both the Gauss-Newton method and the Levenberg-Marquardt method are implemented in the Optimization Toolbox. Details of the implementations are discussed below:

- Gauss-Newton Implementation
- Levenberg-Marquardt Implementation

Gauss-Newton Implementation

The Gauss-Newton method is implemented using similar polynomial line search strategies discussed for unconstrained optimization. In solving the linear least squares problem (Prob. 2.18), exacerbation of the conditioning of the equations is avoided by using the QR decomposition of $J(x_k)$ and applying the decomposition to $F(x_k)$ (using the MATLAB `\` operator). This is in contrast to inverting the explicit matrix, $J(x_k)^T J(x_k)$, which can cause unnecessary errors to occur.

Robustness measures are included in the method. These measures consist of changing the algorithm to the Levenberg-Marquardt method when either the step length goes below a threshold value (in this implementation $1e-15$) or when the condition number of $J(x_k)$ is below $1e-10$. The condition number is a ratio of the largest singular value to the smallest.

Levenberg-Marquardt Implementation

The main difficulty in the implementation of the Levenberg-Marquardt method is an effective strategy for controlling the size of λ_k at each iteration so that it is efficient for a broad spectrum of problems. The method used in this implementation is to estimate the relative nonlinearity of $f(x)$ using a linear predicted sum of squares $f_p(x_k)$ and a cubically interpolated estimate of the minimum $f_k(x_*)$. In this way the size of λ_k is determined at each iteration.

The linear predicted sum of squares is calculated as

$$\begin{aligned} F_p(x_k) &= J(x_{k-1})d_{k-1} + F(x_{k-1}) \\ f_p(x_k) &= F_p(x_k)^T F_p(x_k) \end{aligned} \tag{2-23}$$

and the term $f_k(x_*)$ is obtained by cubically interpolating the points $f(x_k)$ and $f(x_{k-1})$. A step length parameter α^* is also obtained from this interpolation, which is the estimated step to the minimum. If $f_p(x_k)$ is greater than $f_k(x_*)$, then λ_k is reduced, otherwise it is increased. The justification for this is that the difference between $f_p(x_k)$ and $f_k(x_*)$ is a measure of the effectiveness of the Gauss-Newton method and the linearity of the problem. This determines whether to use a direction approaching the steepest descent direction or the Gauss-Newton direction. The formulas for the reduction and increase in λ_k , which have been developed through consideration of a large number of test problems, are shown in Figure 2-5, Updating λ_k below.

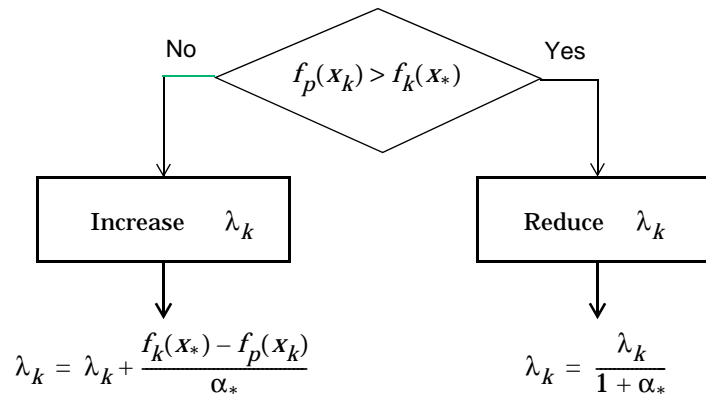


Figure 2-5: Updating λ_k

Following the update of λ_k , a solution of Eq. 2-22 is used to obtain a search direction, d_k . A step length of unity is then taken in the direction d_k , which is followed by a line search procedure similar to that discussed for the unconstrained implementation. The line search procedure ensures that $f(x_{k+1}) < f(x_k)$ at each major iteration and the method is therefore a descent method.

The implementation has been successfully tested on a large number of nonlinear problems. It has proved to be more robust than the Gauss-Newton method and iteratively more efficient than an unconstrained method. The Levenberg-Marquardt algorithm is the default method used by `lsqnonlin`. The Gauss-Newton method can be selected by setting the options parameter `LevenbergMarquardt` to 'off'.

Constrained Optimization

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, which 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 Kuhn-Tucker (KT) equations. The KT 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 KT equations are both necessary and sufficient for a global solution point.

Referring to GP (Eq. 2-1), the Kuhn-Tucker equations can be stated as

$$\begin{aligned} \nabla f(x^*) + \sum_{i=1}^m \lambda_i^* \cdot \nabla G_i(x^*) &= 0 \\ \lambda_i^* \cdot G_i(x^*) &= 0 \quad i = 1, \dots, m \\ \lambda_i^* &\geq 0 \quad i = m_e + 1, \dots, m \end{aligned} \tag{2-24}$$

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. Since 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 zero. This is stated implicitly in the last two equations of Eq. 2-24.

The solution of the KT equations forms the basis to many nonlinear programming algorithms. These algorithms attempt to compute directly the Lagrange multipliers. Constrained quasi-Newton methods guarantee superlinear convergence by accumulating second order information regarding the KT 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).

This section continues with discussions of the following methods for constrained optimization:

- Sequential Quadratic Programming (SQP)
- A Quadratic Programming (QP) Subproblem

“SQP Implementation” on page 2-30 provides implementation details for the Hessian matrix update, quadratic programming problem solution, and line search and merit function calculation phases of the SQP algorithm used in `fmincon`, `fminimax`, `fgoalattain`, and `fseminf`.

Sequential Quadratic Programming (SQP)

SQP methods represent state-of-the-art in nonlinear programming methods. Schittowski [34], 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 ([31],[32]), 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 [33], and Schittowski [23]. The general method, however, is stated here.

Given the problem description in GP (Eq. 2.1) the principal idea is the formulation of a QP subproblem based on a quadratic approximation of the Lagrangian function.

$$L(x, \lambda) = f(x) + \sum_{i=1}^m \lambda_i \cdot g_i(x) \quad (2-25)$$

Here Eq. 2.1 is simplified by assuming that bound constraints have been expressed as inequality constraints. The QP subproblem is obtained by linearizing the nonlinear constraints.

Quadratic Programming (QP) Subproblem

$$\begin{aligned}
 & \underset{d \in \mathfrak{R}^n}{\text{minimize}} && \frac{1}{2} d^T H_k d + \nabla f(x_k)^T d \\
 & && \nabla g_i(x_k)^T d + g_i(x_k) = 0 \quad i = 1, \dots, m_e \\
 & && \nabla g_i(x_k)^T d + g_i(x_k) \leq 0 \quad i = m_e + 1, \dots, m
 \end{aligned} \tag{2-26}$$

This subproblem can be solved using any QP algorithm (see, for instance, “Quadratic Programming Solution” on page 2-31). 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 2-30). The matrix H_k is a positive definite approximation of the Hessian matrix of the Lagrangian function (Eq. 2-25). H_k can be updated by any of the quasi-Newton methods, although the BFGS method (see “Updating the Hessian Matrix” on page 2-30) 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 well-informed decisions regarding directions of search and step length.

Consider Rosenbrock’s function (Eq. 2-2) with an additional nonlinear inequality constraint, $g(x)$

$$x_1^2 + x_2^2 - 1.5 \leq 0 \tag{2-27}$$

This was solved by an SQP implementation in 96 iterations compared to 140 for the unconstrained case. Figure 2-6, SQP Method on Nonlinear Linearly Constrained Rosenbrock’s Function (Eq. 2-2) shows the path to the solution point $x = [0.9072, 0.8228]$ starting at $x = [-1.9, 2]$.




Figure 2-6: SQP Method on Nonlinear Linearly Constrained Rosenbrock's Function (Eq.2-2)

SQP Implementation

The MATLAB SQP implementation consists of three main stages, which are discussed briefly in the following subsections:

- Updating of the Hessian matrix of the Lagrangian function
- Quadratic programming problem solution
- Line search and merit function calculation

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 $\lambda_i (i = 1, \dots, m)$ is an estimate of the Lagrange multipliers.

$$H_{k+1} = H_k + \frac{q_k q_k^T}{q_k^T s_k} - \frac{H_k^T H_k}{s_k^T H_k s_k} \quad \text{where} \quad (2-28)$$

$$s_k = x_{k+1} - x_k$$

$$q_k = \nabla f(x_{k+1}) + \sum_{i=1}^n \lambda_i \cdot \nabla g_i(x_{k+1}) - \left(\nabla f(x_k) + \sum_{i=1}^n \lambda_i \cdot \nabla g_i(x_k) \right)$$

Powell [32] recommends keeping the Hessian positive definite even though it may 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 $q_k \cdot s_k$ is repeatedly halved. This procedure is continued until $q_k^T s_k$ is greater than or equal to $1e-5$. If after this procedure, $q_k^T s_k$ is still not positive, q_k is modified by adding a vector v multiplied by a constant scalar w , that is,

$$q_k = q_k + wv \quad (2-29)$$

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 w is systematically increased until $q_k^T s_k$ becomes positive.

The functions `fmincon`, `fminimax`, `fgoalattain`, and `fseminf` all use SQP. If the options parameter `Display` is set to `'iter'`, 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 procedure described above 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 may 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 is solved of the form where A_i refers to the i th row of the m -by- n matrix A .

$$\begin{aligned} \text{minimize}_{d \in \mathfrak{R}^n} \quad & q(d) = \frac{1}{2} d^T H d + c^T d \\ & A_i d = b_i \quad i = 1, \dots, m_e \\ & A_i d \leq b_i \quad i = m_e + 1, \dots, m \end{aligned} \tag{2-30}$$

The method used in the Optimization Toolbox 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 is maintained, \bar{A}_k , which is an estimate of the active constraints (i.e., which 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

$$Z_k = Q[:, l+1:m] \tag{2-31}$$

where

$$Q^T \bar{A}_k^T = \begin{bmatrix} R \\ 0 \end{bmatrix}$$

Having found Z_k , 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 we view our quadratic as a function of p , by substituting for \hat{d}_k , we have

$$q(p) = \frac{1}{2} p^T Z_k^T H Z_k p + c^T Z_k p \tag{2-32}$$

Differentiating this with respect to p yields

$$\nabla q(p) = Z_k^T H Z_k p + Z_k^T c \quad (2-33)$$

$\nabla q(p)$ 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

$$Z_k^T H Z_k p = -Z_k^T c \quad (2-34)$$

A step is then taken of the form

$$x_{k+1} = x_k + \alpha \hat{d}_k \quad \text{where } \hat{d}_k = Z_k^T p \quad (2-35)$$

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 (Eq. 2.31). 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 iterate. The distance to the constraint boundaries in any direction \hat{d}_k is given by

$$\alpha = \min_i \left\{ \frac{-(A_i x_k - b_i)}{A_i \hat{d}_k} \right\} \quad (i = 1, \dots, m) \quad (2-36)$$

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

$$\bar{A}_k^T \lambda_k = c \quad (2-37)$$

If all elements of λ_k are positive, x_k is the optimal solution of QP (Eq. 2.31). However, if any component of λ_k is negative, and it 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 a point can be found by solving the linear programming problem

$$\begin{aligned}
 & \underset{\gamma \in \mathfrak{R}, x \in \mathfrak{R}^n}{\text{minimize}} && \gamma \\
 & A_i x = b_i && i = 1, \dots, m_e \\
 & A_i x - \gamma \leq b_i && i = m_e + 1, \dots, m
 \end{aligned} \tag{2-38}$$

The notation A_i indicates the i th row of the matrix A . A feasible point (if one exists) to Eq. 2.38 can be found by setting x to a value that satisfies the equality constraints. This can be achieved by solving an under- or over-determined 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.

The above QP algorithm is modified 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).

$$\hat{d}_k = -Z_k Z_k^T g_k \tag{2-39}$$

If a feasible point is found using the above 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

$$H \hat{d}_1 = -g_k \tag{2-40}$$

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

$$x_{k+1} = x_k + \alpha d_k \quad (2-41)$$

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 [32] of the form below has been used in this implementation.

$$\Psi(x) = f(x) + \sum_{i=1}^{m_e} r_i \cdot g_i(x) + \sum_{i=m_e+1}^m r_i \cdot \max\{0, g_i(x)\} \quad (2-42)$$

Powell recommends setting the penalty parameter

$$r_i = (r_{k+1})_i = \max_i \left\{ \lambda_i, \frac{1}{2}((r_k)_i + \lambda_i) \right\}, \quad i = 1, \dots, m \quad (2-43)$$

This allows positive contribution from constraints that are inactive in the QP solution but were recently active. In this implementation, initially the penalty parameter r_i is set to

$$r_i = \frac{\|\nabla f(x)\|}{\|\nabla g_i(x)\|} \quad (2-44)$$

where $\|\ \|$ 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.

Multiobjective Optimization

The rigidity of the mathematical problem posed by the general optimization formulation given in GP (Eq. 2-1) is often remote from that of a practical design problem. Rarely does a single objective with several hard constraints adequately represent the problem being faced. More often there is a vector of objectives $F(x) = \{F_1(x), F_2(x), \dots, F_m(x)\}$ that must be traded off in some way. The relative importance of these objectives is not generally known until the system's best capabilities are determined and trade-offs between the objectives fully understood. As the number of objectives increases, trade-offs are likely to become complex and less easily quantified. There is much reliance on the intuition of the designer and his or her ability to express preferences throughout the optimization cycle. Thus, requirements for a multiobjective design strategy are to enable a natural problem formulation to be expressed, yet be able to solve the problem and enter preferences into a numerically tractable and realistic design problem.

This section includes:

- An introduction to multiobjective optimization, which looks at a number of alternative methods
- A discussion of the Goal Attainment method, which can be posed as a nonlinear programming problem
- Algorithm improvements to the SQP method, for use with the Goal Attainment method.

Introduction

Multiobjective optimization is concerned with the minimization of a vector of objectives $F(x)$ that may be the subject of a number of constraints or bounds.

$$\begin{aligned}
 &\text{minimize } F(x) \\
 &x \in \mathfrak{R}^n \\
 &G_i(x) = 0 \quad i = 1, \dots, m_e \\
 &G_i(x) \leq 0 \quad i = m_e + 1, \dots, m \\
 &x_l \leq x \leq x_u
 \end{aligned} \tag{2-45}$$

Note that, because $F(x)$ is a vector, if any of the components of $F(x)$ are competing, there is no unique solution to this problem. Instead, the concept of

noninferior [37] (also called Pareto optimality [4],[5]) must be used to characterize the objectives. A noninferior solution is one in which an improvement in one objective requires a degradation of another. To define this concept more precisely, consider a feasible region, Ω , in the parameter space $x \in \mathfrak{R}^n$ that satisfies all the constraints, i.e.,

$$\Omega = \{x \in \mathfrak{R}^n\} \quad (2-46)$$

subject to

$$\begin{aligned} g_i(x) &= 0 & i &= 1, \dots, m_e \\ g_i(x) &\leq 0 & i &= m_e + 1, \dots, m \\ x_l &\leq x \leq x_u \end{aligned}$$

This allows us to define the corresponding feasible region for the objective function space Λ .

$$\Lambda = \{y \in \mathfrak{R}^m\} \text{ where } y = F(x) \text{ subject to } x \in \Omega. \quad (2-47)$$

The performance vector, $F(x)$, maps parameter space into objective function space as is represented for a two-dimensional case in Figure 2-7, Mapping from Parameter Space into Objective Function Space, below.




Figure 2-7: Mapping from Parameter Space into Objective Function Space

A noninferior solution point can now be defined.

Definition: A point $x^* \in \Omega$ is a noninferior solution if for some neighborhood of x^* there does not exist a Δx such that $(x^* + \Delta x) \in \Omega$ and

$$\begin{aligned} F_i(x^* + \Delta x) &\leq F_i(x^*) & i = 1, \dots, m \\ F_j(x^* + \Delta x) &< F_j(x^*) & \text{for some } j \end{aligned} \tag{2-48}$$

In the two-dimensional representation of Figure 2-8, Set of Noninferior Solutions, the set of noninferior solutions lies on the curve between C and D . Points A and B represent specific noninferior points.




Figure 2-8: Set of Noninferior Solutions

A and B are clearly noninferior solution points because an improvement in one objective, F_1 , requires a degradation in the other objective, F_2 , i.e., $F_{1B} < F_{1A}$, $F_{2B} > F_{2A}$.

Since any point in Ω that is not a noninferior point represents a point in which improvement can be attained in all the objectives, it is clear that such a point is of no value. Multiobjective optimization is, therefore, concerned with the generation and selection of noninferior solution points. The techniques for multiobjective optimization are wide and varied and all the methods cannot be covered within the scope of this toolbox. However, some of the techniques are described below.

Weighted Sum Strategy

The weighted sum strategy converts the multiobjective problem of minimizing the vector $F(x)$ into a scalar problem by constructing a weighted sum of all the objectives.

$$\underset{x \in \Omega}{\text{minimize}} \quad f(x) = \sum_{i=1}^m w_i \cdot F_i(x)^2 \quad (2-49)$$

The problem can then be optimized using a standard unconstrained optimization algorithm. The problem here is in attaching weighting coefficients to each of the objectives. The weighting coefficients do not necessarily correspond directly to the relative importance of the objectives or allow trade-offs between the objectives to be expressed. Further, the noninferior solution boundary may be nonconcurrent so that certain solutions are not accessible.

This can be illustrated geometrically. Consider the two-objective case in Figure 2-9, Geometrical Representation of the Weighted Sum Method. In the objective function space a line, L , $w^T F(x) = c$ is drawn. The minimization of Eq. 2-49 can be interpreted as finding the value of c for which L just touches the boundary of Λ as it proceeds outwards from the origin. Selection of weights w_1 and w_2 , therefore, defines the slope of L , which in turn leads to the solution point where L touches the boundary of Λ .




Figure 2-9: Geometrical Representation of the Weighted Sum Method

The aforementioned convexity problem arises when the lower boundary of Λ is nonconvex as shown in Figure 2-10, Nonconvex Solution Boundary. In this case the set of noninferior solutions between A and B is not available.




Figure 2-10: Nonconvex Solution Boundary

ε -Constraint Method

A procedure that overcomes some of the convexity problems of the weighted sum technique is the ε -constraint method. This involves minimizing a primary objective, F_p , and expressing the other objectives in the form of inequality constraints

$$\text{minimize}_{x \in \Omega} F_p(x) \quad (2-50)$$

$$\text{subject to } F_i(x) \leq \varepsilon_i \quad i = 1, \dots, m \quad i \neq p$$

Figure 2-11, Geometrical Representation of ε -Constraint Method shows a two-dimensional representation of the ε -constraint method for a two-objective problem.




Figure 2-11: Geometrical Representation of ε -Constraint Method

This approach is able to identify a number of noninferior solutions on a nonconvex boundary that are not obtainable using the weighted sum technique, for example, at the solution point $F_1 = F_{1s}$ and $F_2 = \varepsilon_2$. A problem with this method is, however, a suitable selection of ε to ensure a feasible solution. A further disadvantage of this approach is that the use of hard constraints is rarely adequate for expressing true design objectives. Similar methods exist, such as that of Waltz [36], which prioritize the objectives. The optimization proceeds with reference to these priorities and allowable bounds of acceptance. The difficulty here is in expressing such information at early stages of the optimization cycle.

In order for the designers' true preferences to be put into a mathematical description, the designers must express a full table of their preferences and satisfaction levels for a range of objective value combinations. A procedure must then be realized that is able to find a solution with reference to this. Such methods have been derived for discrete functions using the branches of statistics known as decision theory and game theory (for a basic introduction, see [24]). Implementation for continuous functions requires suitable discretization strategies and complex solution methods. Since it is rare for the designer to know such detailed information, this method is deemed impractical for most practical design problems. It is, however, seen as a possible area for further research.

What is required is a formulation that is simple to express, retains the designers preferences, and is numerically tractable.

Goal Attainment Method

The method described here is the Goal Attainment method of Gembicki [16]. This involves expressing a set of design goals, $F^* = \{F_1^*, F_2^*, \dots, F_m^*\}$, which is 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 over-achieved enabling the designer to be relatively imprecise about initial design goals. The relative degree of under- or over-achievement 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 following formulation.

$$\begin{aligned} & \text{minimize } \gamma \\ & \gamma \in \mathfrak{R}, x \in \Omega \end{aligned} \tag{2-51}$$

$$\text{such that } F_i(x) - w_i \gamma \leq F_i^* \quad i = 1, \dots, m$$

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 trade-offs between the objectives. For instance, setting the weighting vector, w , equal to the initial goals indicates that the same percentage under- or over-attainment of the goals, F^* , is achieved. Hard constraints can be incorporated 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 Goal Attainment method in control system design can be found in Fleming ([10],[11]).

The Goal Attainment method is represented geometrically in Figure 2-12, Geometrical Representation of Goal Attainment Method, for the two-dimensional problem.

$$\begin{aligned} \text{minimize } \gamma \quad & \text{subject to} \quad F_1(x) - w_1\gamma \leq F_1^* \\ \gamma, x \in \Omega \quad & F_2(x) - w_2\gamma \leq F_2^* \end{aligned}$$




Figure 2-12: Geometrical Representation of 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 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, Schittowski [34]). In Goal Attainment programming there may be a more appropriate merit function, which can be achieved by posing Eq. 2-51 as the minimax problem

$$\text{minimize}_{x \in \mathfrak{R}^n} \max_i \{\Lambda_i\} \quad (2-52)$$

$$\text{where } \Lambda_i = \frac{F_i(x) - F_i^*}{w_i} \quad i = 1, \dots, m$$

Following the argument of Brayton et al. [2] for minimax optimization using SQP, using the merit function of Eq. 2-42 for the Goal Attainment problem of Eq. 2-52, gives

$$\psi(x, \gamma) = \gamma + \sum_{i=1}^m r_i \cdot \max \{0, F_i(x) - w_i \gamma - F_i^*\} \quad (2-53)$$

When the merit function of Eq. 2-53 is used as the basis of a line search procedure, then, although $\psi(x, \gamma)$ may decrease for a step in a given search direction, the function $\max \Lambda_i$ may 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, $\psi(x, \gamma)$ may 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.,

$$\psi(x) = \max_i \Lambda_i \quad (2-54)$$

A problem in the Goal Attainment method is that it is common to use a weighting coefficient equal to zero to incorporate hard constraints. The merit function of Eq. 2-54 then becomes infinite for arbitrary violations of the constraints. To overcome this problem while still retaining the features of Eq. 2-54 the merit function is combined with that of Eq. 2-43 giving the following.

$$\psi(x) = \sum_{i=1}^m \begin{cases} r_i \cdot \max \{0, F_i(x) - w_i \gamma - F_i^*\} & \text{if } w_i = 0 \\ \max_i \Lambda_i \quad i = 1, \dots, m & \text{otherwise} \end{cases} \quad (2-55)$$

Another feature that can be exploited in SQP is the objective function γ . From the KT equations (Eq. 2-24) 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 . By initializing H as the identity matrix, this property does not appear. 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 2-31.

The above modifications have been implemented in `fgoalattain` and have been found to make the method more robust. However, due to 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 (Eq. 2-42).

Selected 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] Censor, Y., "Pareto Optimality in Multiobjective Problems," *Appl. Math. Optimiz.*, Vol. 4, pp. 41-59, 1977.
- [5] Da Cunha, N.O. and E. Polak, "Constrained Minimization Under Vector-valued Criteria in Finite Dimensional Spaces," *J. Math. Anal. Appl.*, Vol. 19, pp. 103-124, 1967.
- [6] Dantzig, G., *Linear Programming and Extensions*, Princeton University Press, Princeton, 1963.
- [7] Dantzig, G., A. Orden, and P. Wolfe, "Generalized Simplex Method for Minimizing a Linear from Under Linear Inequality Constraints," *Pacific J. Math.* Vol. 5, pp. 183-195.
- [8] Davidon, W.C., "Variable Metric Method for Minimization," *A.E.C. Research and Development Report*, ANL-5990, 1959.
- [9] Dennis, J.E., Jr., "Nonlinear Least Squares," *State of the Art in Numerical Analysis* ed. D. Jacobs, Academic Press, pp. 269-312, 1977.
- [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 Porg. 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," *Vol. 1, Unconstrained Optimization*, and *Vol. 2, Constrained Optimization*, John Wiley and Sons., 1980.
- [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. Dissertation, 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*, Academic Press, London, 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. Schittowski, "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] 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.
- [32] 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.
- [33] 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.
- [34] Schittowski, K., "NLQPL: A FORTRAN-Subroutine Solving Constrained Nonlinear Programming Problems," *Annals of Operations Research*, Vol. 5, pp. 485-500, 1985.
- [35] Shanno, D.F., "Conditioning of Quasi-Newton Methods for Function Minimization," *Mathematics of Computing*, Vol. 24, pp. 647-656, 1970.
- [36] Waltz, F.M., "An Engineering Approach: Hierarchical Optimization Criteria," *IEEE Trans.*, Vol. AC-12, pp. 179-180, April, 1967.
- [37] Zadeh, L.A., "Optimality and Nonscalar-valued Performance Criteria," *IEEE Trans. Automat. Contr.*, Vol. AC-8, p. 1, 1963.

Large-Scale Algorithms

Overview	3-2
Trust Region Methods for Nonlinear Minimization . . .	3-3
Preconditioned Conjugate Gradients	3-6
Linearly Constrained Problems	3-8
Linear Equality Constraints	3-8
Box Constraints	3-8
Nonlinear Least Squares	3-11
Quadratic Programming	3-12
Linear Least Squares	3-13
Large-Scale Linear Programming	3-14
Main Algorithm	3-14
Preprocessing	3-17
Selected Bibliography	3-18

Overview

To solve large-scale optimization problems, special techniques are needed. This chapter describes the techniques used in the Optimization Toolbox for the large-scale methods. It is organized into these major topics.

Trust Region Methods for Nonlinear Minimization

Introduces the trust region idea, and describes the use of trust region ideas for unconstrained nonlinear minimization.

Preconditioned Conjugate Gradients

Presents an algorithm that uses Preconditioned Conjugate Gradients (PCG) for solving large symmetric positive definite systems of linear equations.

Linearly Constrained Problems

Discusses the solution of linear equality constrained and box constrained minimization problems.

Nonlinear Least Squares

Describes the solution of nonlinear least squares problems.

Quadratic Programming

Describes the solution of minimization problems with quadratic objective functions.

Linear Least Squares

Describes the solution of linear least squares problems.

Large-Scale Linear Programming

Describes the use of LIPSOL (Linear Interior Point Solver) for the solution of large-scale linear programming problems.

Selected Bibliography

Lists published materials that support concepts implemented in the large-scale algorithms.

Trust Region Methods for Nonlinear Minimization

Many of the methods used in the Optimization Toolbox are based on the trust region idea, 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 we are at a point x in n -space and we 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,

$$\min_s \{q(s) \quad s \in N\} \quad (3-1)$$

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. In this section we focus on the unconstrained problem; additional complications due to the presence of constraints on the variables are discussed in later sections.

In the standard trust region method ([8]), 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

$$\min \left\{ \frac{1}{2} s^T H s + s^T g \quad \text{such that} \quad \|D s\| \leq \Delta \right\} \quad (3-2)$$

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 Eq. 3-2 (see [8]); 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 Eq. 3-2, 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 Eq. 3-2, have been proposed in the literature ([2],[10]). The approximation approach we follow in the Optimization Toolbox is to restrict the trust region subproblem to a two-dimensional subspace S ([1],[2]). Once the subspace S has been computed, the work to solve Eq. 3-2 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 toolbox assigns $S = \langle s_1, s_2 \rangle$ 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

$$H \cdot s_2 = -g \tag{3-3}$$

or a direction of negative curvature

$$s_2^T \cdot H \cdot s_2 < 0 \tag{3-4}$$

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 framework for the Optimization Toolbox approach to unconstrained minimization using trust region ideas is now easy to describe:

- Formulate the two-dimensional trust region subproblem
- Solve Eq. 3-2 to determine the trial step s
- If $f(x+s) \leq f(x)$ then $x = x+s$
- 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 [6],[9] for a discussion of this aspect.

The Optimization Toolbox treats 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 Gradients

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 perform matrix-vector products of the form $H \cdot v$ 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.

Algorithm

The Optimization Toolbox uses this PCG algorithm, which it refers to as Algorithm PCG.

```
% Initializations
r = -g; p = zeros(n,1);
% Precondition
z = M\r; inner1 = r'*z; inner2 = 0; d = z;
% Conjugate gradient iteration
for k = 1:kmax
    if k > 1
        beta = inner1/inner2;
        d = z + beta*d;
    end
    w = H*d; denom = d'*w;
    if denom <= 0
        p = d/norm(d); % Direction of negative/zero curvature
        break % Exit if zero/negative curvature detected
    else
        alpha = inner1/denom;
        p = p + alpha*d;
        r = r - alpha*w;
    end
    z = M\r;
    if norm(z) <= tol % Exit if Hp=-g solved within tolerance
        break
    end
    inner2 = inner1;
    inner1 = r'*z;
end
```

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 $H p = -g$. In either case p is used to help define the two-dimensional subspace used in the trust region approach discussed in the section “Trust Region Methods for Nonlinear Minimization” on page 3-3.

Linearly Constrained Problems

Linear constraints complicate the situation described for unconstrained minimization. However, the underlying ideas described above can be carried through in a clean and efficient way. The large-scale methods in the Optimization Toolbox generate strictly feasible iterates:

- A projection technique is used for linear equality constraints.
- Reflections are used with simple box constraints.

Linear Equality Constraints

The general linear equality constrained minimization problem can be written

$$\min \{f(x) \text{ such that } Ax = b\} \quad (3-5)$$

where A is an m -by- n matrix ($m \leq n$). The Optimization Toolbox preprocesses A to remove strict linear dependencies using a technique based on the LU-factorization of A^T [6]. Here we will assume A is of rank m .

Our method to solve Eq. 3-5 differs from our 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 [6], 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

$$\begin{bmatrix} C & \tilde{A}^T \\ \tilde{A} & 0 \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix} \quad (3-6)$$

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., $C \approx H$. See [6] for more details.

Box Constraints

The box constrained problem is of the form

$$\min \{f(x) \text{ such that } l \leq x \leq u\} \quad (3-7)$$

where l is a vector of lower bounds, and u is a vector of upper bounds. Some (or all) of the components of l may be equal to $-\infty$ and some (or all) of the components of u may 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 Eq. 3-7

$$(D(x))^{-2}g = 0 \quad (3-8)$$

where

$$D(x) = \text{diag} \left(|v_k|^{\frac{1}{2}} \right)$$

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 Eq. 3-8 is not differentiable everywhere; nondifferentiability occurs when $v_i = 0$. Hence we avoid such points by maintaining strict feasibility, i.e., restricting $l < x < u$.

The scaled modified Newton step s_k^N for Eq. 3-8 is defined as the solution to the linear system

$$\hat{M}Ds^N = -\hat{g} \quad (3-9)$$

where

$$\hat{g} = D^{-1}g = \text{diag} \left(|v|^{\frac{1}{2}} \right)g \quad (3-10)$$

and

$$\hat{M} = D^{-1}HD^{-1} + \text{diag}(g)J^v \quad (3-11)$$

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 may not be differentiable. We define $J_{ii}^v = 0$ 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 $|v_i| \cdot g_i$ 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 $p_i^R = -p_i$.

Nonlinear Least Squares

An important special case for $f(x)$ is the nonlinear least squares problem

$$f(x) = \frac{1}{2} \sum_i f_i^2(x) = \frac{1}{2} \|F(x)\|_2^2 \quad (3-12)$$

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 3-3. 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

$$\min \|Js + F\|_2^2 \quad (3-13)$$

(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.

Quadratic Programming

In this case the function $f(x)$ is the quadratic equation

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

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 [5] for details of the line search.

Linear Least Squares

In this case the function $f(x)$ to be solved is

$$f(x) = \frac{1}{2} \|Cx + d\|_2^2$$

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.,

$$C^T Cx = -C^T d$$

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 [5] 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.

Large-Scale Linear Programming

Linear programming is defined as

$$\min f^T x \quad \text{such that} \quad \begin{pmatrix} A_{eq} \cdot x = b_{eq} \\ A_{ineq} \cdot x \leq b_{ineq} \\ l \leq x \leq u \end{pmatrix} \quad (3-14)$$

The large-scale method is based on LIPSOL ([11]), which is a variant of Mehrotra's predictor-corrector algorithm ([7]), a primal-dual interior-point method.

This section continues with descriptions of:

- The main algorithm
- Preprocessing steps

Main Algorithm

The algorithm begins by applying a series of preprocessing steps (see "Preprocessing" on page 3-17). After preprocessing, the problem has the form

$$\min f^T x \quad \text{such that} \quad \begin{pmatrix} A \cdot x = b \\ 0 \leq x \leq u \end{pmatrix} \quad (3-15)$$

The upper bounds constraints are implicitly included in the constraint matrix A with the addition of primal slack variables s , Eq. 3-15 becomes

$$\min f^T x \quad \text{such that} \quad \begin{pmatrix} A \cdot x = b \\ x + s = u \\ x \geq 0, s \geq 0 \end{pmatrix} \quad (3-16)$$

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

$$\max b^T y - u^T w \quad \text{such that} \quad \begin{matrix} A^T \cdot y - w + z = f \\ z \geq 0, w \geq 0 \end{matrix} \quad (3-17)$$

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 Eq. 3-16 and the dual Eq. 3-17, are

$$F(x, y, z, s, w) = \begin{pmatrix} A \cdot x - b \\ x + s - u \\ A^T \cdot y - w + z - f \\ x_i z_i \\ s_i w_i \end{pmatrix} = 0 \quad (3-18)$$

$$x \geq 0, z \geq 0, s \geq 0, w \geq 0$$

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 Eq. 3-18, 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 Eq. 3-16.)

The algorithm is a variant of the predictor-corrector algorithm proposed by Mehrotra. Consider an iterate $v = [x; y; z; s; w]$, where $[x; z; s; w] > 0$. We first compute the so-called *prediction* direction

$$\Delta v_p = -(F^T(v))^{-1} F(v)$$

which is the Newton direction; then the so-called *corrector* direction

$$\Delta v_c = -(F^T(v))^{-1} (F(v + \Delta v_p)) - \mu \hat{e}$$

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. We combine the two directions 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 steps above, the algorithm computes a (sparse) direct factorization on a modification of the Cholesky factors of $A \cdot A^T$. If A has dense columns, it instead uses the Sherman-Morrison formula, and if that solution is not adequate (the residual is too large), it uses preconditioned conjugate gradients to find a solution.

The algorithm then repeats these steps 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 Eq. 3-18.

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, and not the original, problem. Thus, if the multipliers are *not* requested, this transformation back will not be computed, and may save some time computationally.

Selected Bibliography

- [1] Branch, M.A., T.F. Coleman, 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.
- [2] 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.
- [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] 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.
- [5] 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.
- [6] Coleman, T.F. and A. Verma, "A Preconditioned Conjugate Gradient Approach to Linear Equality Constrained Minimization," submitted to *Computational Optimization and Applications*.
- [7] Mehrotra, S., "On the Implementation of a Primal-Dual Interior Point Method," *SIAM Journal on Optimization*, Vol. 2, pp. 575-601, 1992.
- [8] 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.
- [9] 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.
- [10] Steihaug, T., "The Conjugate Gradient Method and Trust Regions in Large Scale Optimization," *SIAM Journal on Numerical Analysis*, Vol. 20, pp. 626-637, 1983.
- [11] 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.

Function Reference

Functions by Category	4-3
Minimization	4-3
Equation Solving	4-3
Least Squares (Curve Fitting)	4-4
Utility	4-4
Demos of Large-Scale Methods	4-4
Demos of Medium-Scale Methods	4-5
Function Arguments	4-6
Optimization Options Parameters	4-11

This chapter provides detailed descriptions of the functions in the Optimization Toolbox. It contains:

- Functions by Category
- Alphabetical List of Functions
- Function Arguments
- Optimization Options Parameters

Functions by Category

This section lists the Optimization Toolbox functions by the type of optimization to which they apply. It also lists demos of large-scale methods and medium-scale methods.

Minimization

Function	Purpose
fgoalattain	Multiobjective goal attainment
fminbnd	Scalar nonlinear minimization with bounds
fmincon	Constrained nonlinear minimization
fminimax	Minimax optimization
fminsearch, fminunc	Unconstrained nonlinear minimization
fseminf	Semi-infinite minimization
linprog	Linear programming
quadprog	Quadratic programming

Equation Solving

Function	Purpose
\	Use \ (left division) to solve linear equations. See the Arithmetic Operators reference page.
fsolve	Nonlinear equation solving
fzero	Scalar nonlinear equation solving

Least Squares (Curve Fitting)

Function	Purpose
<code>\</code>	Use <code>\</code> (left division) for linear least squares with no constraints. See the Arithmetic Operators reference page.
<code>lsqlin</code>	Constrained linear least squares
<code>lsqcurvefit</code>	Nonlinear curve fitting
<code>lsqnonlin</code>	Nonlinear least squares
<code>lsqnonneg</code>	Nonnegative linear least squares

Utility

Function	Purpose
<code>optimset</code> , <code>optimget</code>	Parameter setting

Demos of Large-Scale Methods

From the MATLAB Help Browser, click the demo name to run the demo. Look for information and additional instructions in the MATLAB command window.

Demo	Purpose
<code>circustent</code>	Quadratic programming to find shape of a circus tent
<code>molecule</code>	Molecule conformation solution using unconstrained nonlinear minimization
<code>optdeblur</code>	Image deblurring using bounded linear least squares

Demos of Medium-Scale Methods

From the MATLAB Help Browser, click the demo name to run the demo. Look for information and additional instructions in the MATLAB command window.

Demo	Purpose
bandemo	Minimization of the banana function
dfildemo	Finite-precision filter design (requires the Signal Processing Toolbox)
goaldemo	Goal attainment example
optdemo	Menu of demo routines
tutdemo	Script for the medium scale algorithms. The script follows the “Tutorial” chapter of the Optimization Toolbox User’s Guide.

Function Arguments

These tables describe the input arguments and output arguments used by Optimization Toolbox functions. They also list the functions that use each argument.

See the individual function reference pages for function-specific information about these arguments.

Table 4-1: 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: $A*x \leq b$.	fgoalattain, fmincon, fminimax, fseminf, 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: $Aeq*x = beq$.	fgoalattain, fmincon, fminimax, fseminf, linprog, lsqlin, quadprog
C, d	The matrix C and vector d are, respectively, the coefficients of the over- or under-determined linear system and the right-hand-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 $x' * H * x + f' * x$.	linprog, quadprog
fun	The function to be optimized. fun is a function or an inline object. See the individual function reference pages for more information on fun.	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fseminf, fsolve, fzero, lsqcurvefit, lsqnonlin
goal	Vector of values that the objectives attempt to attain. The vector is the same length as the number of objectives.	fgoalattain

Table 4-1: Input Arguments (Continued)

Argument	Description	Used by Functions
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. Unbounded variables may also be specified using -Inf (for lower bounds) or Inf (for upper bounds). For example, if $lb(i) = -Inf$ then 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. See the individual reference pages for more information on nonlcon.	fgoalattain, fmincon, fminimax
ntheta	The number of semi-infinite constraints.	fseminf
options	An optimization options parameter structure that defines parameters used by the optimization functions. For information about the parameters, see Table 4-3, Optimization Options Parameters or the individual function reference pages.	All functions
P1, P2, ...	Additional arguments to be passed to fun, nonlcon (if it exists), and seminfcon (if it exists), when the optimization function calls the functions fun, nonlcon, or seminfcon using these calls: $f = \text{feval}(\text{fun}, x, P1, P2, \dots)$ $[c, \text{ceq}] = \text{feval}(\text{nonlcon}, x, P1, P2, \dots)$ $[c, \text{ceq}, K1, K2, \dots, Kn, s] = \dots$ $\text{feval}(\text{seminfcon}, x, s, P1, P2, \dots)$	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fseminf, fsolve, fzero, lsqcurvefit, lsqnonlin

Table 4-1: Input Arguments (Continued)

Argument	Description	Used by Functions
	Using this feature, the same <code>fun</code> (or <code>nonlcon</code> or <code>seminfcon</code>) can solve a number of similar problems with different parameters, avoiding the need to use global variables.	
<code>seminfcon</code>	The function that computes the nonlinear inequality and equality constraints <i>and</i> the semi-infinite constraints. <code>seminfcon</code> is the name of an M-file or MEX-file. See the individual function reference pages for <code>fseminf</code> more information on <code>seminfcon</code> .	<code>fseminf</code>
<code>weight</code>	A weighting vector to control the relative under-attainment or over-attainment of the objectives.	<code>fgoalattain</code>
<code>xdata</code> , <code>ydata</code>	The input data <code>xdata</code> and the observed output data <code>ydata</code> that is to be fit to an equation.	<code>lsqcurvefit</code>
<code>x0</code>	Starting point (a scalar, vector or matrix). (For <code>fzero</code> , <code>x0</code> can also be a two-element vector representing an interval that is known to contain a zero.)	All functions except <code>fminbnd</code>
<code>x1</code> , <code>x2</code>	The interval over which the function is minimized.	<code>fminbnd</code>

Table 4-2: 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>	The exit condition. For the meaning of a particular value, see the function reference pages.	All functions

Table 4-2: Output Arguments (Continued)

Argument	Description	Used by Functions
fval	The value of the objective function fun at the solution x.	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fseminf, fsolve, fzero, linprog, quadprog
grad	The value of the gradient of fun at the solution x. If fun does not compute the gradient, grad is a finite-differencing approximation of the gradient.	fmincon, fminunc
hessian	The value of the Hessian of fun at the solution x. For large-scale methods, if fun does not compute the Hessian, hessian is a finite-differencing approximation of the Hessian. For medium-scale methods, hessian is the value of the Quasi-Newton approximation to the Hessian at the solution x.	fmincon, fminunc
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. 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

Table 4-2: Output Arguments (Continued)

Argument	Description	Used by Functions
output	An output structure that contains information about the results of the optimization. For structure field names, see individual function descriptions.	All functions
residual	The value of the residual at the solution x .	lsqcurvefit, lsqin, lsqnonlin, lsqnonneg
resnorm	The value of the squared 2-norm of the residual at the solution x .	lsqcurvefit, lsqin, lsqnonlin, lsqnonneg
x	The solution found by the optimization function. If <code>exitflag > 0</code> then x is a solution; otherwise, x is the value the optimization routine was at when it terminated prematurely.	All functions

Optimization Options Parameters

This table describes fields in the optimization parameters structure, `options`. The column labeled L, M, B indicates whether the parameter applies to large-scale methods, medium scale methods, or both:

- L – Large-scale methods only
- M – Medium-scale methods only
- B – Both large- and medium-scale methods

See the Optimization Toolbox `optimset` reference page, the MATLAB `optimset` reference page, and the individual function reference pages for information about parameter values and defaults.

Note Links in this table are to Optimization Toolbox functions. These are links to the corresponding MATLAB optimization functions: `fminbnd`, `fminsearch`, `fzero`, `lsqnonneg`, `optimget`, `optimset`.

Table 4-3: Optimization Options Parameters

Parameter Name	Description	L, M, B	Used by Functions
<code>DerivativeCheck</code>	Compare user-supplied analytic derivatives (gradients or Jacobian) to finite differencing derivatives.	M	<code>fgoalattain</code> , <code>fmincon</code> , <code>fminimax</code> , <code>fminunc</code> , <code>fseminf</code> , <code>fsolve</code> , <code>lsqcurvefit</code> , <code>lsqnonlin</code>
<code>Diagnostics</code>	Print diagnostic information about the function to be minimized or solved.	B	All but <code>fminbnd</code> , <code>fminsearch</code> , <code>fzero</code> , and <code>lsqnonneg</code>
<code>DiffMaxChange</code>	Maximum change in variables for finite difference derivatives.	M	<code>fgoalattain</code> , <code>fmincon</code> , <code>fminimax</code> , <code>fminunc</code> , <code>fseminf</code> , <code>fsolve</code> , <code>lsqcurvefit</code> , <code>lsqnonlin</code>

Table 4-3: Optimization Options Parameters (Continued)

Parameter Name	Description	L, M, B	Used by Functions
DiffMinChange	Minimum change in variables for finite difference derivatives.	M	fgoalattain, fmincon, fminimax, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin
Display	Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' displays just the final output; 'notify' displays output only if function does not converge.	B	All. See the individual function reference pages for the values that apply.
GoalsExactAchieve	Number of goals to achieve exactly (do not over- or underachieve).	M	fgoalattain
GradConstr	Gradients for the nonlinear constraints defined by the user.	M	fgoalattain, fmincon, fminimax
GradObj	Gradient(s) for the objective function(s) defined by the user.	B	fgoalattain, fmincon, fminimax, fminunc, fseminf
Hessian	If 'on', function uses user-defined Hessian, or Hessian information (when using HessMult), for the objective function. If 'off', function approximates the Hessian using finite differences.	L	fmincon, fminunc
HessMult	Hessian multiply function defined by the user.	L	fmincon, fminunc, quadprog

Table 4-3: Optimization Options Parameters (Continued)

Parameter Name	Description	L, M, B	Used by Functions
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 x0, the starting point.	L	fmincon, fminunc
HessUpdate	Quasi-Newton updating scheme.	M	fminunc
Jacobian	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.	B	fsolve, lsqcurvefit, lsqnonlin
JacobMult	Jacobian multiply function defined by the user.	L	fsolve, lsqcurvefit, lsqlin, lsqnonlin
JacobPattern	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.	L	fsolve, lsqcurvefit, lsqnonlin
LargeScale	Use large-scale algorithm if possible.	B	fmincon, fminunc, fsolve, linprog, lsqcurvefit, lsqlin, lsqnonlin, quadprog

Table 4-3: Optimization Options Parameters (Continued)

Parameter Name	Description	L, M, B	Used by Functions
LevenbergMarquardt	Chooses Levenberg-Marquardt over Gauss-Newton algorithm.	M	fsolve, lsqcurvefit, lsqnonlin
LineSearchType	Line search algorithm choice.	M	fminunc, fsolve, lsqcurvefit, lsqnonlin
MaxFunEvals	Maximum number of function evaluations allowed.	B	fgoalattain, fminbnd, fmincon, fminimax, fminsearch, fminunc, fseminf, fsolve, lsqcurvefit, lsqnonlin
MaxIter	Maximum number of iterations allowed.	B	All but fzero and lsqnonneg
MaxPCGIter	Maximum number of PCG iterations allowed.	L	fmincon, fminunc, fsolve, lsqcurvefit, lsqlin, lsqnonlin, quadprog
MeritFunction	Use goal attainment/minimax merit function (multiobjective) vs. fmincon (single objective).	M	fgoalattain, fminimax
MinAbsMax	Number of $F(x)$ to minimize the worst case absolute values	M	fminimax
PrecondBandWidth	Upper bandwidth of preconditioner for PCG.	L	fmincon, fminunc, fsolve, lsqcurvefit, lsqlin, lsqnonlin, quadprog
TolCon	Termination tolerance on the constraint violation.	B	fgoalattain, fmincon, fminimax, fseminf

Table 4-3: Optimization Options Parameters (Continued)

Parameter Name	Description	L, M, B	Used by Functions
TolFun	Termination tolerance on the function value.	B	fgoalattain, fmincon, fminimax, fminsearch, fminunc, fseminf, fsolve, linprog (large-scale only), lsqcurvefit, lsqin (large-scale only), lsqnonlin, quadprog (large-scale only)
TolPCG	Termination tolerance on the PCG iteration.	L	fmincon, fminunc, fsolve, lsqcurvefit, lsqin, lsqnonlin, quadprog
TolX	Termination tolerance on x .	B	All functions except the medium-scale algorithms for linprog, lsqin, and quadprog
TypicalX	Typical x values. The length of the vector is equal to the number of elements in x_0 , the starting point.	L	fmincon, fminunc, fsolve, lsqcurvefit, lsqin, lsqnonlin, quadprog

fgoalattain

Purpose

Solve multiobjective goal attainment problem

$$\begin{array}{ll} \text{minimize}_{x,\gamma} \gamma & \text{such that} \\ & F(x) - \text{weight} \cdot \gamma \leq \text{goal} \\ & c(x) \leq 0 \\ & \text{ceq}(x) = 0 \\ & A \cdot x \leq b \\ & A_{\text{eq}} \cdot x = \text{beq} \\ & lb \leq x \leq ub \end{array}$$

where x , weight , goal , b , beq , lb , and ub are vectors, A and A_{eq} are matrices, and $c(x)$, $\text{ceq}(x)$, and $F(x)$ are functions that return vectors. $F(x)$, $c(x)$, and $\text{ceq}(x)$ can be nonlinear functions.

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(fun,x0,goal,weight,A,b,Aeq,beq,...
    lb,ub,nonlcon,options,P1,P2,...)
[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.

`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 x_0 , with weight specified by `weight`.

`x = fgoalattain(fun,x0,goal,weight,A,b)` solves the goal attainment problem subject to the linear inequalities $A \cdot x \leq b$.

`x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq)` solves the goal attainment problem subject to the linear equalities $Aeq \cdot x = beq$ 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, `x`, so that the solution is always in the range $lb \leq x \leq ub$.

`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 parameters specified in the structure `options`.

`x = fgoalattain(fun,x0,goal,weight,A,b,Aeq,beq,lb,ub,nonlcon,... options,P1,P2,...)` passes the problem-dependent parameters `P1`, `P2`, etc., directly to the functions `fun` and `nonlcon`. Pass empty matrices as placeholders for `A`, `b`, `Aeq`, `beq`, `lb`, `ub`, `nonlcon`, and `options` if these arguments are not needed.

`[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`.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `fgoalattain`. This section provides function-specific details for `fun`, `goal`, `nonlcon`, `options`, and `weight`:

`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.

```
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 an inline object.

```
x = fgoalattain(inline('sin(x.*x)'),x0,goal,weight);
```

To make an objective function as near as possible to a goal value, (i.e., neither greater than nor less than) set `options.GoalsExactAchieve` 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* `options.GradObj` is 'on', as set by

```
options = optimset('GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `G`, a matrix, at `x`. Note that by checking the value of `nargout` the function can avoid computing `G` when `fun` 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 nargout > 1   % Two output arguments
    G = ...       % Gradients evaluated at x
end
```


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 x_0 , 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 $F(j)$).

`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* `options.GradConstr` is 'on', as set by

```
options = optimset('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 $ceq(x)$. Note that by checking the value of `nargout` the function can avoid computing `GC` and `GCEq` when `nonlcon` is called with only two output arguments (in the case where the optimization algorithm only needs the values of c and ceq but not `GC` and `GCEq`).

```
function [c,ceq,GC,GCEq] = mycon(x)
c = ...           % Nonlinear inequalities at x
ceq = ...         % Nonlinear equalities at x
if nargin > 2     % Nonlcon called with 4 outputs
    GC = ...      % Gradients of the inequalities
    GCEq = ...    % Gradients of the equalities
end
```

If `nonlcon` returns a vector `c` of m components and `x` has length n , where n is the length of `x0`, 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 $ceq(j)$).

options “Options” on page 4-22 provides the function-specific details for the options parameters.

weight A weighting vector to control the relative under-attainment or over-attainment of the objectives in `fgoalattain`. When the values of `goal` are *all nonzero*, to ensure the same percentage of under- or over-attainment 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 `weight=abs(goal)` when any of the `goal` values are zero causes that goal constraint to be treated like a hard constraint rather than as a goal constraint.

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` parameter and put that objective as the first element of the vector returned by `fun` (see the description of `fun` and options above).

Output Arguments. Table 4-2, Output Arguments, 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. If <code>attainfactor</code> is negative, the goals have been over-achieved; if <code>attainfactor</code> is positive, the goals have been under-achieved.
<code>exitflag</code>	Describes the exit condition: <ul style="list-style-type: none"> > 0 The function converged to a solution x. 0 The maximum number of function evaluations or iterations was exceeded. < 0 The function did not converge to a solution.
<code>lambda</code>	Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields of the structure are: <ul style="list-style-type: none"> <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:

iterations	Number of iterations taken.
funcCount	Number of function evaluations.
algorithm	Algorithm used.

Options

Optimization options parameters used by `fgoalattain`. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information:

DerivativeCheck	Compare user-supplied derivatives (gradients of objective or constraints) to finite-differencing derivatives.
Diagnostics	Print diagnostic information about the function to be minimized or solved.
DiffMaxChange	Maximum change in variables for finite-difference gradients.
DiffMinChange	Minimum change in variables for finite-difference gradients.
Display	Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.
GoalsExactAchieve	Specifies the number of goals to “just” achieve, that is, do not try to over- or underachieve.
GradConstr	Gradient for the constraints defined by user. See the description of <code>nonlcon</code> above to see how to define the gradient in <code>nonlcon</code>
GradObj	Gradient for the objective function defined by user. See the description of <code>fun</code> above to see how to define the gradient in <code>fun</code> . The gradient <i>must</i> be provided to use the large-scale method. It is optional for the medium-scale method.
MaxFunEvals	Maximum number of function evaluations allowed.
MaxIter	Maximum number of iterations allowed.

MeritFunction	Use goal attainment/minimax merit function if set to 'multiobj'. Use fmincon merit function if set to 'singleobj'.
TolCon	Termination tolerance on the constraint violation.
TolFun	Termination tolerance on the function value.
TolX	Termination tolerance on x.

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 over-attainment in the active objectives at the solution, the weighting matrix, `weight`, is set to `abs(goal)`.

Starting with a controller, $K = [-1, -1; -1, -1]$, first write an M-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 = optimset('Display','iter'); % Set display parameter
[K,fval,attainfactor] = fgoalattain(@eigfun,K0,...
    goal,weight,[],[],[],[],lb,ub,[],options,A,B,C)
```

This example can be run by using the demonstration script `goaldemo`. After about 12 iterations, a solution is

```
Active constraints:
    1
    2
    4
    9
   10
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 over-achieved 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 over-attainment is met exactly. Three of the lower bound constraints are also active.

In the above 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, set `options.GoalsExactAchieve` to the number of objectives for which this is required.

Consider the above 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 `options.GoalsExactAchieve` set to 3.

```
options = optimset('GoalsExactAchieve',3);
[K,fval,attainfactor] = fgoalattain(...
    @eigfun,K0,goal,weight,[],[],[],[],lb,ub,[],options,A,B,C)
```

After about seven iterations, a solution is

```
K =
    -1.5954    1.2040
    -0.4201   -2.9046

fval =
    -5.0000
    -3.0000
    -1.0000

attainfactor =
    1.0859e-20
```

In this case the optimizer has tried to match the objectives to the goals. The attainment factor (of $1.0859e-20$) 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 since 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.

Algorithm

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 fully in “Standard Algorithms” on page 2-1.

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 unknown whether the objectives will reach the goals (under attainment) or be minimized less than the goals (over attainment). A weighting vector, *weight*, controls the relative under-attainment or over-attainment of the objectives.

fgoalattain uses a Sequential Quadratic Programming (SQP) method, which is described fully in “Standard Algorithms” on page 2-1. 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], [6]. The line search is terminated when either merit function shows improvement. A modified Hessian, which takes advantage of special structure of this problem, is also used (see [1] and [[4]). A full description of the modifications used is found in “Goal Attainment Method” on page 2-42 in “Introduction to Algorithms.” Setting `options.MeritFunction = 'singleobj'` uses the merit function and Hessian used in `fmincon`.

`attainfactor` contains the value of γ at the solution. A negative value of γ indicates over-attainment in the goals.

See also “SQP Implementation” on page 2-30 in “Introduction to Algorithms” for more details on the algorithm used and the types of procedures printed under the Procedures heading for the `options.Display = 'iter'` setting.

Limitations

The objectives must be continuous. fgoalattain may give only local solutions.

See Also

@(function_handle), `fmincon`, `fminimax`, `optimset`

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*, Springer Verlag, Vol. 630, 1978.

fminbnd

Purpose

Find the minimum of a function of one variable on a fixed interval

$$\min_x f(x) \quad \text{such that} \quad x_1 < x < x_2$$

where 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(fun,x1,x2,options,P1,P2,...)
[x,fval] = fminbnd(...)
[x,fval,exitflag] = fminbnd(...)
[x,fval,exitflag,output] = fminbnd(...)
```

Description

`fminbnd` finds the minimum of a function of one variable within a fixed interval.

`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 $x_1 < x < x_2$.

`x = fminbnd(fun,x1,x2,options)` minimizes with the optimization parameters specified in the structure `options`.

`x = fminbnd(fun,x1,x2,options,P1,P2,...)` provides for additional arguments, `P1`, `P2`, etc., which are passed to the objective function, `fun`. Use `options=[]` as a placeholder if no options are set.

`[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.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `fminbnd`. This section provides function-specific details for `fun` and `options`:

fun The function to be minimized. `fun` is 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.

```
x = fminbnd(@myfun,x0)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...           % Compute function value at x.
```

`fun` can also be an inline object.

```
x = fminbnd(inline('sin(x*x)'),x0);
```

options “Options” on page 4-29 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `fminbnd`. This section provides function-specific details for `exitflag` and output:

exitflag Describes the exit condition:

> 0	The function converged to a solution <code>x</code> .
0	The maximum number of function evaluations or iterations was exceeded.
< 0	The function did not converge to a solution.

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>algorithm</code>	Algorithm used.

Options

Optimization options parameters used by `fminbnd`. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information:

fminbnd

Display	Level of display. '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.
MaxFunEvals	Maximum number of function evaluations allowed.
MaxIter	Maximum number of iterations allowed.
TolX	Termination tolerance on x.

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 an M-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 = f(x)
```

$$y = -1$$

Algorithm `fminbnd` is an M-file. The algorithm is based on Golden Section search and parabolic interpolation. A Fortran program implementing the same algorithm is given in [1].

Limitations The function to be minimized must be continuous. `fminbnd` may 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.

See Also `@(function_handle)`, `fminsearch`, `fmincon`, `fminunc`, `optimset`, `inline`

References [1] Forsythe, G.E., M.A. Malcolm, and C.B. Moler, *Computer Methods for Mathematical Computations*, Prentice Hall, 1976.

fmincon

Purpose

Find the minimum of a constrained nonlinear multivariable function

$$\min_x f(x) \text{ subject to}$$

$$c(x) \leq 0$$

$$ceq(x) = 0$$

$$A \cdot x \leq b$$

$$Aeq \cdot x = beq$$

$$lb \leq x \leq ub$$

where x , b , beq , lb , and ub 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.

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(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options,P1,P2, ...)
[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` finds the 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*.

`x = fmincon(fun,x0,A,b)` starts at x_0 and finds a minimum x to the function described in `fun` subject to the linear inequalities $A*x \leq b$. x_0 can be a scalar, vector, or matrix.

`x = fmincon(fun,x0,A,b,Aeq,beq)` minimizes `fun` subject to the linear equalities $Aeq*x = beq$ as well as $A*x \leq b$. Set $A=[]$ and $b=[]$ if no inequalities exist.

`x = fmincon(fun,x0,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 = fmincon(fun,x0,A,b,Aeq,beq,lb,ub,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 $ceq(x) = 0$. Set `lb=[]` and/or `ub=[]` if no bounds exist.

`x = fmincon(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)` minimizes with the optimization parameters specified in the structure `options`.

`x = fmincon(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options,P1,P2,...)` passes the problem-dependent parameters `P1`, `P2`, etc., directly to the functions `fun` and `nonlcon`. Pass empty matrices as placeholders for `A`, `b`, `Aeq`, `beq`, `lb`, `ub`, `nonlcon`, and `options` if these arguments are not needed.

`[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 of `fun` at the solution x .

Arguments

Input Arguments. Table 4-1, **Input Arguments**, contains general descriptions of arguments passed in to `fmincon`. This “Arguments” section provides function-specific details for `fun`, `nonlcon`, and `options`:

`fun` The function to be minimized. `fun` is 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.

```
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 an inline object.

```
x = fmincon(inline('norm(x)^2'),x0,A,b);
```

If the gradient of `fun` can also be computed *and* `options.GradObj` is `'on'`, as set by

```
options = optimset('GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `g`, a vector, at `x`. Note that by checking the value of `nargout` the function can avoid computing `g` when `fun` 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 value at x
if nargout > 1   % fun called with two output arguments
    g = ...       % Compute the gradient evaluated at x
end
```

The gradient consists of the partial derivatives 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* `options.Hessian` is `'on'`, i.e., `options = optimset('Hessian','on')`, then the function `fun` must return the Hessian value `H`, a symmetric matrix, at `x` in a third output argument. Note that by checking the value of `nargout` we can avoid computing `H` when `fun` is called with only one or two output arguments (in the case where the optimization algorithm only needs the values of `f` and `g` but not `H`).


```

function [f,g,H] = myfun(x)
f = ...      % Compute the objective function value at x
if nargin > 1 % fun called with two output arguments
    g = ...  % Gradient of the function evaluated at x
    if nargin > 2
        H = ... % Hessian evaluated at x
    end
end

```

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.

`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 = 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* `options.GradConstr` is 'on', as set by

```
options = optimset('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 $ceq(x)$. Note that by checking the value of `nargout` the function can avoid computing GC and $GCeq$ when `nonlcon` is called with only two output arguments (in the case where the optimization algorithm only needs the values of c and ceq but not GC and $GCeq$).

```
function [c,ceq,GC,GCEq] = mycon(x)
c = ...           % Nonlinear inequalities at x
ceq = ...         % Nonlinear equalities at x
if nargin > 2     % nonlcon called with 4 outputs
    GC = ...      % Gradients of the inequalities
    GCEq = ...    % Gradients of the equalities
end
```

If `nonlcon` returns a vector `c` of m components and `x` has length n , where n is the length of `x0`, 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 $ceq(j)$).

`options` “Options” on page 4-37 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `fmincon`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

`exitflag` Describes the exit condition:

> 0	The function converged to a solution x .
0	The maximum number of function evaluations or iterations was exceeded.
< 0	The function did not converge to a solution.

`lambda` 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

	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.
	algorithm	Algorithm used.
	cgiterations	Number of PCG iterations (large-scale algorithm only).
	stepsize	Final step size taken (medium-scale algorithm only).
	firstorderopt	Measure of first-order optimality (large-scale algorithm only). For large-scale 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 3-8, and g is the gradient. For large-scale problems with only linear equalities, the first-order optimality is the infinity norm of the <i>projected</i> gradient (i.e. the gradient projected onto the nullspace of Aeq).

Options

Optimization options parameters used by `fmincon`. Some parameters apply to all algorithms, some are only relevant when using the large-scale algorithm, and others are only relevant when using the medium-scale algorithm. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information.

We start by describing the `LargeScale` option since it states a *preference* for which algorithm to use. It is only a preference since certain conditions must be met to use the large-scale algorithm. For `fmincon`, you must provide the gradient (see the description of `fun` above to see how) or else the medium-scale algorithm is used:

LargeScale Use large-scale algorithm if possible when set to 'on'.
Use medium-scale algorithm when set to 'off'.

Medium-Scale and Large-Scale Algorithms. These parameters are used by both the medium-scale and large-scale algorithms:

Diagnostics Print diagnostic information about the function to be minimized.

Display Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.

GradObj Gradient for the objective function defined by user. See the description of `fun` above to see how to define the gradient in `fun`. You must provide the gradient to use the large-scale method. It is optional for the medium-scale method.

MaxFunEvals Maximum number of function evaluations allowed.

MaxIter Maximum number of iterations allowed.

TolFun Termination tolerance on the function value.

TolCon Termination tolerance on the constraint violation.

TolX Termination tolerance on x .

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

Hessian If 'on', `fmincon` uses a user-defined Hessian (defined in `fun`), or Hessian information (when using `HessMult`), for the objective function. If 'off', `fmincon` 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

$$W = \text{hmfun}(\text{Hinfo}, Y, p1, p2, \dots)$$

where H_{info} and the additional parameters p_1, p_2, \dots contain the matrices used to compute $H*Y$. The first argument must be the same as the third argument returned by the objective function `fun`.

```
[f,g,Hinfo] = fun(x,p1,p2,...)
```

The parameters p_1, p_2, \dots are the same additional parameters that are passed to `fmincon` (and to `fun`).

```
fmincon(fun,...,options,p1,p2,...)
```

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 H_{info} to compute the preconditioner.

Note 'Hessian' must be set to 'on' for H_{info} to be passed from `fun` to `hmfun`.

See “Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints” on page 1-48 for an example.

HessPattern

Sparsity pattern of the Hessian for finite-differencing. If it is not convenient to compute the sparse Hessian matrix H in `fun`, the large-scale method in `fmincon` can approximate H via sparse finite-differences (of the gradient) provided the *sparsity structure* of H — i.e., locations of the nonzeros — is supplied as the value for `HessPattern`. In the worst case, if the structure is unknown, you can set `HessPattern` to be a dense matrix and a full finite-difference approximation is computed at each iteration (this is the default). This can be very expensive for large problems so it is usually worth the effort to determine the sparsity structure.

MaxPCGIter

Maximum number of PCG (preconditioned conjugate gradient) iterations (see the *Algorithm* section below).

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.
TolPCG	Termination tolerance on the PCG iteration.
TypicalX	Typical x values.

Medium-Scale Algorithm Only. These parameters are used only by the medium-scale algorithm:

DerivativeCheck	Compare user-supplied derivatives (gradients of the objective and constraints) to finite-differencing derivatives.
DiffMaxChange	Maximum change in variables for finite-difference gradients.
DiffMinChange	Minimum change in variables for finite-difference gradients.

Examples

Find values of x that minimize $f(x) = -x_1 x_2 x_3$, starting at the point $x = [10; 10; 10]$ and subject to the constraints

$$0 \leq x_1 + 2x_2 + 2x_3 \leq 72$$

First, write an M-file that returns a scalar value f of the function evaluated at x .

```
function f = myfun(x)
f = -x(1) * x(2) * x(3);
```

Then rewrite the constraints as both less than or equal to a constant,

$$-x_1 - 2x_2 - 2x_3 \leq 0$$

$$x_1 + 2x_2 + 2x_3 \leq 72$$

Since both constraints are linear, formulate them as the matrix inequality $A \cdot x \leq b$ where

$$A = \begin{bmatrix} -1 & -2 & -2 \\ 1 & 2 & 2 \end{bmatrix} \quad b = \begin{bmatrix} 0 \\ 72 \end{bmatrix}$$

Next, 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)
```

After 66 function evaluations, the solution is

```
x =
    24.0000
    12.0000
    12.0000
```

where the function value is

```
fval =
   -3.4560e+03
```

and linear inequality constraints evaluate to be ≤ 0

```
A*x-b=
   -72
     0
```

Notes

Large-Scale Optimization. To use the large-scale method, the gradient must be provided in `fun` (and `options.GradObj` set to `'on'`). A warning is given if no gradient is provided and `options.LargeScale` is not `'off'`. The function `fmincon` permits $g(x)$ to be an approximate gradient but this option is not recommended; the numerical behavior of most optimization codes is considerably more robust when the true gradient is used.

The large-scale method in `fmincon` is 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` parameter in `options`), then `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, then `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.

Several aspects of linearly constrained minimization should be noted:

- A dense (or fairly dense) column of matrix `Aeq` can result in considerable fill and computational cost.
- `fmincon` removes (numerically) linearly dependent rows in `Aeq`; however, this process involves repeated matrix factorizations and therefore can be costly if there are many dependencies.
- Each iteration involves a sparse least squares solve with matrix

$$\overline{Aeq} = Aeq^T R^{-T}$$

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} .

Medium-Scale Optimization. Better numerical results are likely if you specify equalities explicitly using `Aeq` and `beq`, instead of implicitly using `lb` and `ub`.

If equality constraints are present and dependent equalities are detected and removed in the quadratic subproblem, 'dependent' is printed under the Procedures heading (when output is asked for using `options.Display = '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' is printed under the Procedures heading.

Algorithm

Large-Scale Optimization. By default `fmincon` will choose the large-scale algorithm *if* the user supplies the gradient in `fun` (and `GradObj` is 'on' in `options`) *and* if *only* upper and lower bounds exist *or* *only* linear equality constraints exist. This algorithm is a subspace trust region method and is based on the interior-reflective Newton method described in [1], [2]. 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 “Large-Scale Algorithms” on page 3-1.

Medium-Scale Optimization. `fmincon` uses a Sequential Quadratic Programming (SQP) method. In this method, a Quadratic Programming (QP) subproblem is solved at each iteration. An estimate of the Hessian of the Lagrangian is updated at each iteration using the BFGS formula (see `fminunc`, references [7], [8]).

A line search is performed using a merit function similar to that proposed by [4], [5], and [6]. The QP subproblem is solved using an active set strategy similar to that described in [3]. A full description of this algorithm is found in “Constrained Optimization” on page 2-26 in “Introduction to Algorithms.”

See also “SQP Implementation” on page 2-30 in “Introduction to Algorithms” for more details on the algorithm used.

Diagnostics

Large-Scale Optimization. The large-scale code does not allow equal upper and lower bounds. For example if `lb(2)==ub(2)`, then `fmincon` gives the error

```
Equal upper and lower bounds not permitted in this large-scale method.
```

```
Use equality constraints and the medium-scale method instead.
```

If you only have equality constraints you can still use the large-scale method. But if you have both equalities and bounds, you must use the medium-scale method.

Limitations

The function to be minimized and the constraints must both be continuous. `fmincon` may only give local solutions.

When the problem is infeasible, `fmincon` attempts to minimize the maximum constraint value.

The objective function and constraint function must be real-valued, that is they cannot return complex values.

Large-Scale Optimization. To use the large-scale algorithm, the user must supply the gradient in `fun` (and `GradObj` must be set 'on' in options), and only upper and lower bounds constraints may be specified, *or only* linear equality constraints must exist and `Aeq` cannot have more rows than columns. `Aeq` is typically sparse. See Table 1-4, Large-Scale Problem Coverage and

Requirements for more information on what problem formulations are covered and what information must be provided.

Currently, if the analytical gradient is provided in `fun`, the options parameter `DerivativeCheck` cannot be used with the large-scale method to compare the analytic gradient to the finite-difference gradient. Instead, use the medium-scale method to check the derivative with options parameter `MaxIter` set to 0 iterations. Then run the problem with the large-scale method.

See Also

@(function_handle), `fminbnd`, `fminsearch`, `fminunc`, `optimset`

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

[4] Han, S.P., "A Globally Convergent Method for Nonlinear Programming," *Journal of Optimization Theory and Applications*, Vol. 22, p. 297, 1977.

[5] 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.

[6] 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.

Purpose

Solve the minimax problem

$$\min_x \max_{\{F_i\}} \{F_i(x)\} \quad \text{such that} \quad \begin{aligned} c(x) &\leq 0 \\ ceq(x) &= 0 \\ A \cdot x &\leq b \\ Aeq \cdot x &= beq \\ lb &\leq x \leq ub \end{aligned}$$

where x , b , beq , lb , and ub are vectors, A and Aeq are matrices, $c(x)$, $ceq(x)$, and $F(x)$ are functions that return vectors. $F(x)$, $c(x)$, and $ceq(x)$ can be nonlinear functions.

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(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options,P1,P2,...)
[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 value of a set of multivariable functions, starting at an initial estimate. The values may be subject to constraints. This is generally referred to as the *minimax* problem.

`x = fminimax(fun,x0)` starts at x_0 and finds a minimax solution x to the functions described in `fun`.

`x = fminimax(fun,x0,A,b)` solves the minimax problem subject to the linear inequalities $A \cdot x \leq b$.

fminimax

`x = fminimax(fun,x,A,b,Aeq,beq)` solves the minimax problem subject to the linear equalities $Aeq \cdot x = beq$ as well. Set $A=[]$ and $b=[]$ if no inequalities exist.

`x = fminimax(fun,x,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 $lb \leq x \leq ub$.

`x = fminimax(fun,x0,A,b,Aeq,beq,lb,ub,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.

`x = fminimax(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)` minimizes with the optimization parameters specified in the structure options.

`x = fminimax(fun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options,P1,P2,...)` passes the problem-dependent parameters $P1$, $P2$, etc., directly to the functions `fun` and `nonlcon`. Pass empty matrices as placeholders for A , b , Aeq , beq , lb , ub , `nonlcon`, and `options` if these arguments are not needed.

`[x,fval] = fminimax(...)` returns the value of the objective function `fun` at the solution x .

`[x,fval,maxfval] = fminimax(...)` returns the maximum function value 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 .

Arguments

Input Arguments. Table 4-1, **Input Arguments**, contains general descriptions of arguments passed in to `fminimax`. This section provides function-specific details for `fun`, `nonlcon`, and `options`:

`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.

```
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 an inline object.

```
x = fminimax(inline('sin(x.*x)'),x0);
```

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 set options.`MinAbsMax` to be the number of such objectives.

If the gradient of the objective function can also be computed *and* options.`GradObj` is 'on', as set by

```
options = optimset('GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `G`, a matrix, at `x`. Note that 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 nargout > 1   % Two output arguments
    G = ...       % Gradients evaluated at x
end
```

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 $F(j)$).

`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* options.`GradConstr` is 'on', as set by

```
options = optimset('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 $ceq(x)$. Note that by checking the value of `nargout` the function can avoid computing `GC` and `GCEq` when `nonlcon` is called with only two output arguments (in the case where the optimization algorithm only needs the values of c and ceq but not `GC` and `GCEq`).

```
function [c,ceq,GC,GCEq] = mycon(x)
c = ...          % Nonlinear inequalities at x
ceq = ...        % Nonlinear equalities at x
if nargout > 2  % nonlcon called with 4 outputs
    GC = ...     % Gradients of the inequalities
    GCEq = ...   % Gradients of the equalities
end
```

If `nonlcon` returns a vector `c` of `m` components and `x` has length `n`, where `n` is the length of `x0`, 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)`).

`options` “Options” on page 4-50 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `fminimax`. This section provides function-specific details for `exitflag`, `lambda`, `maxfval`, and `output`:

`exitflag` Describes the exit condition:

- > 0 The function converged to a solution `x`.
- 0 The maximum number of function evaluations or iterations was exceeded.
- < 0 The function did not converge to a solution.

`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

`maxfval` Maximum of the function values evaluated at the solution `x`, that is, `maxfval = max{fun(x)}`.

fminimax

output Structure containing information about the optimization. The fields of the structure are:

- iterations Number of iterations taken.
- funcCount Number of function evaluations.
- algorithm Algorithm used.

Options

Optimization options parameters used by `fminimax`. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information:

- DerivativeCheck Compare user-supplied derivatives (gradients of the objective or constraints) to finite-differencing derivatives.
- Diagnostics Print diagnostic information about the function to be minimized or solved.
- DiffMaxChange Maximum change in variables for finite-difference gradients.
- DiffMinChange Minimum change in variables for finite-difference gradients.
- Display Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.
- GradConstr Gradient for the constraints defined by user. See the description of `nonlcon` above to see how to define the gradient in `nonlcon`.
- GradObj Gradient for the objective function defined by user. See the description of `fun` above to see how to define the gradient in `fun`. The gradient *must* be provided to use the large-scale method. It is optional for the medium-scale method.
- MaxFunEvals Maximum number of function evaluations allowed.
- MaxIter Maximum number of iterations allowed.

MeritFunction	Use goal attainment/minimax merit function if set to 'multiobj'. Use fmincon merit function if set to 'singleobj'.
MinAbsMax	Number of $F(x)$ to minimize the worst case absolute values.
TolCon	Termination tolerance on the constraint violation.
TolFun	Termination tolerance on the function value.
TolX	Termination tolerance on x.

Examples

Find values of x that minimize the maximum value of

$$\left[f_1(x), f_2(x), f_3(x), f_4(x), f_5(x) \right]$$

where

$$f_1(x) = 2x_1^2 + x_2^2 - 48x_1 - 40x_2 + 304$$

$$f_2(x) = -x_2^2 - 3x_2^2$$

$$f_3(x) = x_1 + 3x_2 - 18$$

$$f_4(x) = -x_1 - x_2$$

$$f_5(x) = x_1 + x_2 - 8.$$

First, write an M-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

fminimax

```
x =
    4.0000
    4.0000
fval =
    0.0000  -64.0000  -2.0000  -8.0000  -0.0000
```

Notes

The number of objectives for which the worst case absolute values of F are minimized is set in `options.MinAbsMax`. Such objectives should be partitioned into the first elements of F .

For example, consider the above problem, which requires finding values of x that minimize the maximum absolute value of

$$\left[f_1(x), f_2(x), f_3(x), f_4(x), f_5(x) \right]$$

This is solved by invoking `fminimax` with the commands

```
x0 = [0.1; 0.1];           % Make a starting guess at the solution
options = optimset('MinAbsMax',5); % Minimize absolute 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
```

If equality constraints are present and dependent equalities are detected and removed in the quadratic subproblem, 'dependent' is printed under the Procedures heading (when the output option is `options.Display='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' is printed under the Procedures heading.

Algorithm

`fminimax` uses a Sequential Quadratic Programming (SQP) method [1]. 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. A modified Hessian that takes advantage of

special structure of this problem is also used. Setting `options.MeritFunction = 'singleobj'` uses the merit function and Hessian used in `fmincon`.

See also “SQP Implementation” on page 2-30 in “Introduction to Algorithms” for more details on the algorithm used and the types of procedures printed under the `Procedures` heading for `options.Display='iter'` setting.

Limitations

The function to be minimized must be continuous. `fminimax` may only give local solutions.

See Also

@(function_handle), `fgoalattain`, `lsqnonlin`, `optimset`

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 Transactions of Circuits and Systems*, Vol. CAS-26, Sept. 1979.

[5] Powell, M.J.D., “A Fast Algorithm for Nonlinear Constrained Optimization Calculations,” *Numerical Analysis*, ed. G.A. Watson, *Lecture Notes in Mathematics*, Springer Verlag, Vol. 630, 1978.

fminsearch

Purpose Find the minimum of an unconstrained multivariable function

$$\min_x f(x)$$

where x is a vector and $f(x)$ is a function that returns a scalar.

Syntax

```
x = fminsearch(fun,x0)
x = fminsearch(fun,x0,options)
x = fminsearch(fun,x0,options,P1,P2,...)
[x,fval] = fminsearch(...)
[x,fval,exitflag] = fminsearch(...)
[x,fval,exitflag,output] = fminsearch(...)
```

Description

`fminsearch` finds the minimum of a scalar function of several variables, starting at an initial estimate. This is generally referred to as *unconstrained nonlinear optimization*.

`x = fminsearch(fun,x0)` starts at the point `x0` and finds a local minimum `x` of the function described in `fun`. `x0` can be a scalar, vector, or matrix.

`x = fminsearch(fun,x0,options)` minimizes with the optimization parameters specified in the structure `options`.

`x = fminsearch(fun,x0,options,P1,P2,...)` passes the problem-dependent parameters `P1`, `P2`, etc., directly to the function `fun`. Use `options = []` as a placeholder if no options are set.

`[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.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `fminsearch`. This section provides function-specific details for `fun` and `options`:

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.

```
x = fminsearch(@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 an inline object.

```
x = fminsearch(inline('norm(x)^2'),x0,A,b);
```

options “Options” on page 4-55 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `fminsearch`. This section provides function-specific details for `exitflag` and output:

exitflag Describes the exit condition:

- > 0 The function converged to a solution `x`.
- 0 The maximum number of function evaluations or iterations was exceeded.
- < 0 The function did not converge to a solution.

output Structure containing information about the optimization. The fields of the structure are:

- `iterations` Number of iterations taken.
- `funcCount` Number of function evaluations.
- `algorithm` Algorithm used.

Options

Optimization options parameters used by `fminsearch`. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information:

fminsearch

Display	Level of display. '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.
MaxFunEvals	Maximum number of function evaluations allowed.
MaxIter	Maximum number of iterations allowed.
TolFun	Termination tolerance on the function value.
TolX	Termination tolerance on x.

Examples

Minimize the one-dimensional function $f(x) = \sin(x) + 3$.

To use an M-file, i.e., `fun = 'myfun'`, create a file `myfun.m`.

```
function f = myfun(x)
f = sin(x) + 3;
```

Then call `fminsearch` to find a minimum of `fun` near 2.

```
x = fminsearch(@myfun,2)
```

To minimize the function $f(x) = \sin(x) + 3$ using an inline object

```
f = inline('sin(x)+3');
x = fminsearch(f,2);
```

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`.

If n is the length of x , a simplex in n -dimensional space is characterized by the $n+1$ distinct vectors that are its vertices. In two-space, a simplex is a triangle; in three-space, it is a pyramid. At each step of the search, a new point in or near the current simplex is generated. The function value at the new point is compared with the function's values at the vertices of the simplex and, usually, one of the vertices is replaced by the new point, giving a new simplex. This step is repeated until the diameter of the simplex is less than the specified tolerance.

`fminsearch` is generally less efficient than `fminunc` for problems of order greater than two. However, when the problem is highly discontinuous, `fminsearch` may be more robust.

Limitations

fminsearch can often handle discontinuity, particularly if it does not occur near the solution. fminsearch may 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.

Note fminsearch is not the preferred choice for solving problems that are sums-of-squares, that is, of the form $\min f(x) = f_1(x)^2 + f_2(x)^2 + f_3(x)^2 + L$. Instead use the `lsqnonlin` function, which has been optimized for problems of this form.

See Also

@(function_handle), fminbnd, fminunc, inline, optimset

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.

fminunc

Purpose Find the minimum of an unconstrained multivariable function

$$\min_x f(x)$$

where x is a vector and $f(x)$ is a function that returns a scalar.

Syntax

```
x = fminunc(fun,x0)
x = fminunc(fun,x0,options)
x = fminunc(fun,x0,options,P1,P2,...)
[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` finds the minimum of a scalar function of several variables, starting at an initial estimate. This is generally referred to as *unconstrained nonlinear optimization*.

`x = fminunc(fun,x0)` starts at the point `x0` and finds 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 parameters specified in the structure options.

`x = fminunc(fun,x0,options,P1,P2,...)` passes the problem-dependent parameters `P1`, `P2`, etc., directly to the function `fun`. Pass an empty matrix for options to use the default values for options.

`[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`.

Arguments

Input Arguments. Table 4-1, **Input Arguments**, contains general descriptions of arguments passed in to `fminunc`. This section provides function-specific details for `fun` and options:

`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.

```
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 an inline object.

```
x = fminunc(inline('norm(x)^2'),x0);
```

If the gradient of `fun` can also be computed *and* options.`GradObj` is 'on', as set by

```
options = optimset('GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `g`, a vector, at `x`. Note that by checking the value of `nargout` the function can avoid computing `g` when `fun` 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 value at x
if nargout > 1   % fun called with 2 output arguments
    g = ...       % Compute the gradient evaluated at x
end
```

The gradient is the partial derivatives $\partial f/\partial x$ 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* options.Hessian is 'on', i.e., options = optimset('Hessian','on'), then the function fun must return the Hessian value H, a symmetric matrix, at x in a third output argument. Note that by checking the value of nargout we can avoid computing H when fun is called with only one or two output arguments (in the case where the optimization algorithm only needs the values of f and g but not H).

```
function [f,g,H] = myfun(x)
f = ... % Compute the objective function value at x
if nargout > 1 % fun called with two output arguments
    g = ... % Gradient of the function evaluated at x
    if nargout > 2
        H = ... % Hessian evaluated at x
    end
end
```

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.

options “Options” on page 4-61 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by fminunc. This section provides function-specific details for exitflag and output:

exitflag Describes the exit condition:

- | | |
|-----|--|
| > 0 | The function converged to a solution x . |
| 0 | The maximum number of function evaluations or iterations was exceeded. |
| < 0 | The function did not converge to a solution. |

output	Structure containing information about the optimization. The fields of the structure are:
iterations	Number of iterations taken.
funcCount	Number of function evaluations.
algorithm	Algorithm used.
cgiterations	Number of PCG iterations (large-scale algorithm only).
stepsize	Final step size taken (medium-scale algorithm only).
firstorderopt	Measure of first-order optimality: the norm of the gradient at the solution x .

Options

fminunc uses these optimization parameters. Some parameters apply to all algorithms, some are only relevant when using the large-scale algorithm, and others are only relevant when using the medium-scale algorithm. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information.

We start by describing the `LargeScale` option since it states a *preference* for which algorithm to use. It is only a preference since certain conditions must be met to use the large-scale algorithm. For `fminunc`, the *gradient must be provided* (see the description of `fun` above to see how) or else the minimum-scale algorithm is used:

<code>LargeScale</code>	Use large-scale algorithm if possible when set to 'on'. Use medium-scale algorithm when set to 'off'.
-------------------------	--

Large-Scale and Medium-Scale Algorithms. These parameters are used by both the large-scale and medium-scale algorithms:

<code>Diagnostics</code>	Print diagnostic information about the function to be minimized.
--------------------------	--

Display	Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.
GradObj	Gradient for the objective function defined by user. See the description of fun above to see how to define the gradient in fun. The gradient <i>must</i> be provided to use the large-scale method. It is optional for the medium-scale method.
MaxFunEvals	Maximum number of function evaluations allowed.
MaxIter	Maximum number of iterations allowed.
TolFun	Termination tolerance on the function value.
TolX	Termination tolerance on x .

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

Hessian If 'on', fminunc uses a user-defined Hessian (defined in fun), or Hessian information (when using HessMult), for the objective function. If 'off', 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

$$W = \text{hmfun}(\text{Hinfo}, Y, p1, p2, \dots)$$

where Hinfo and the additional parameters $p1, p2, \dots$ contain the matrices used to compute $H*Y$.

The first argument must be the same as the third argument returned by the objective function fun.

$$[f, g, \text{Hinfo}] = \text{fun}(x, p1, p2, \dots)$$

The parameters p_1, p_2, \dots are the same additional parameters that are passed to `fminunc` (and to `fun`).

```
fminunc(fun, ..., options, p1, p2, ...)
```

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.

Note 'Hessian' must be set to 'on' for `Hinfo` to be passed from `fun` to `hmfun`.

See “Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints” on page 1-48 for an example.

HessPattern

Sparsity pattern of the Hessian for finite-differencing. If it is not convenient to compute the sparse Hessian matrix H in `fun`, the large-scale method in `fminunc` can approximate H via sparse finite-differences (of the gradient) provided the *sparsity structure* of H — i.e., locations of the nonzeros — is supplied as the value for `HessPattern`. In the worst case, if the structure is unknown, you can set `HessPattern` to be a dense matrix and a full finite-difference approximation is computed at each iteration (this is the default). This can be very expensive for large problems so it is usually worth the effort to determine the sparsity structure.

MaxPCGIter

Maximum number of PCG (preconditioned conjugate gradient) iterations (see the *Algorithm* section below).

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.

TolPCG Termination tolerance on the PCG iteration.
TypicalX Typical x values.

Medium-Scale Algorithm Only. These parameters are used only by the medium-scale algorithm:

DerivativeCheck Compare user-supplied derivatives (gradient) to finite-differencing derivatives.
DiffMaxChange Maximum change in variables for finite-difference gradients.
DiffMinChange Minimum change in variables for finite-difference gradients.
LineSearchType Line search algorithm choice.

Examples

Minimize the function $f(x) = 3x_1^2 + 2x_1x_2 + x_2^2$.

To use an M-file, 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 couple of iterations, the solution, `x`, and the value of the function at `x`, `fval`, are returned.

```
x =
  1.0e-008 *
   -0.7914    0.2260
fval =
  1.5722e-016
```

To minimize this function with the gradient provided, modify the M-file `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

```

and indicate the gradient value is available by creating an optimization options structure with `options.GradObj` set to 'on' using `optimset`.

```

options = optimset('GradObj','on');
x0 = [1,1];
[x,fval] = fminunc(@myfun,x0,options)

```

After several iterations the solution `x` and `fval`, the value of the function at `x`, are returned.

```

x =
    1.0e-015 *
   -0.6661      0
fval2 =
    1.3312e-030

```

To minimize the function $f(x) = \sin(x) + 3$ using an inline object

```

f = inline('sin(x)+3');
x = fminunc(f,4)

```

which returns a solution

```

x =
    4.7124

```

Notes

`fminunc` is not the preferred choice for solving problems that are sums-of-squares, that is, of the form

$$\min f(x) = f_1(x)^2 + f_2(x)^2 + f_3(x)^2 + L$$

Instead use the `lsqnonlin` function, which has been optimized for problems of this form.

To use the large-scale method, the gradient must be provided in `fun` (and `options.GradObj` set to 'on'). A warning is given if no gradient is provided and `options.LargeScale` is not 'off'.

Algorithms

Large-Scale Optimization. By default `fminunc` chooses the large-scale algorithm if the user supplies the gradient in `fun` (and `GradObj` is 'on' in options). This algorithm is a subspace trust region method and is based on the interior-reflective Newton method described in [2],[3]. 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 3-3, and “Preconditioned Conjugate Gradients” on page 3-6 in the “Large-Scale Algorithms” section.

Medium-Scale Optimization. `fminunc` with options.LargeScale set to 'off' uses the BFGS Quasi-Newton method with a mixed quadratic and cubic line search procedure. This quasi-Newton method uses the BFGS ([1],[5],[8],[9]) formula for updating the approximation of the Hessian matrix. The DFP ([4],[6],[7]) formula, which approximates the inverse Hessian matrix, is selected by setting options.HessUpdate to 'dfp' (and options.LargeScale to 'off'). A steepest descent method is selected by setting options.HessUpdate to 'steepdesc' (and options.LargeScale to 'off'), although this is not recommended.

The default line search algorithm, i.e., when options.LineSearchType is set to 'quadcubic', is a safeguarded mixed quadratic and cubic polynomial interpolation and extrapolation method. A safeguarded cubic polynomial method can be selected by setting options.LineSearchType to 'cubicpoly'. This second method generally requires fewer function evaluations but more gradient evaluations. Thus, if gradients are being supplied and can be calculated inexpensively, the cubic polynomial line search method is preferable. A full description of the algorithms is given in “Standard Algorithms” on page 2-1.

Limitations

The function to be minimized must be continuous.`fminunc` may 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.

Large-Scale Optimization. To use the large-scale algorithm, the user must supply the gradient in `fun` (and `GradObj` must be set 'on' in options). See Table 1-4, Large-Scale Problem Coverage and Requirements for more information on what problem formulations are covered and what information must be provided.

Currently, if the analytical gradient is provided in `fun`, the options parameter `DerivativeCheck` cannot be used with the large-scale method to compare the analytic gradient to the finite-difference gradient. Instead, use the medium-scale method to check the derivative with options parameter `MaxIter` set to 0 iterations. Then run the problem again with the large-scale method.

See Also

@(function_handle), `fminsearch`, `inline`, `optimset`

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.

fseminf

Purpose

Find minimum of a semi-infinitely constrained multivariable nonlinear function

$$\begin{array}{ll} \min_x f(x) & \text{subject to} \\ & c(x) \leq 0, \\ & ceq(x) = 0 \\ & A \cdot x \leq b \\ & Aeq \cdot x = beq \\ & lb \leq x \leq ub \\ & K_1(x, w_1) \leq 0 \\ & K_2(x, w_2) \leq 0 \\ & \dots \\ & K_n(x, w_n) \leq 0 \end{array}$$

where x , b , beq , lb , and ub 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.

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(fun,x0,ntheta,seminfcon,A,b,Aeq,beq,...
            lb,ub,options,P1,P2,...)
[x,fval] = fseminf(...)
[x,fval,exitflag] = fseminf(...)
[x,fval,exitflag,output] = fseminf(...)
[x,fval,exitflag,output,lambda] = fseminf(...)
```

Description

`fseminf` finds the 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 $w_i \in \mathfrak{R}^1$ (or $w_i \in \mathfrak{R}^2$). Since it is impossible to calculate all possible values of $K_f(x, w_i)$, a region must be chosen for w_i over which to calculate an appropriately sampled set of values.

`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 $A*x \leq b$.

`x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq)` minimizes subject to the linear equalities $Aeq*x = beq$ 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, `x`, so that the solution is always in the range $lb \leq x \leq ub$.

`x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq,lb,ub,options)` minimizes with the optimization parameters specified in the structure `options`.

`x = fseminf(fun,x0,ntheta,seminfcon,A,b,Aeq,beq,lb,ub,options,P1,P2,...)` passes the problem-dependent parameters `P1`, `P2`, etc., directly to the functions `fun` and `seminfcon`. Pass empty matrices as placeholders for `A`, `b`, `Aeq`, `beq`, `lb`, `ub`, and `options` if these arguments are not needed.

`[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`.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `fseminf`. This section provides function-specific details for `fun`, `ntheta`, `options`, and `seminfcon`:

`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.

```
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 an inline object.

```
fun = inline('sin(x'*x)');
```

If the gradient of `fun` can also be computed *and* `options.GradObj` is 'on', as set by

```
options = optimset('GradObj','on')
```

then the function `fun` must return, in the second output argument, the gradient value `g`, a vector, at `x`. Note that by checking the value of `nargout` the function can avoid computing `g` when `fun` 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 value at x
if nargout > 1   % fun called with 2 output arguments
    g = ...       % Compute the gradient evaluated at x
end
```

The gradient is the partial derivatives 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`.

`ntheta` The number of semi-infinite constraints.

`options` “Options” on page 4-72 provides the function-specific details for the options parameters.

`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 may or may not be used. Return `[]` for `c` and `ceq` if no such constraints exist.

The vectors or matrices, `K1`, `K2`, ..., `Kntheta`, contain the semi-infinite constraints evaluated for a sampled set of values for the independent variables, `w1`, `w2`, ..., `wntheta`, respectively. The two column matrix, `S`, contains a recommended sampling interval for values of `w1`, `w2`, ..., `wntheta`, which are used to evaluate `K1`, `K2`, ..., `Kntheta`. The *i*th row of `S` contains the recommended sampling interval for evaluating `Ki`. When `Ki` is a vector, use only `S(i,1)` (the second column can be all zeros). When `Ki` is a matrix, `S(i,2)` is used for the sampling of the rows in `Ki`, `S(i,1)` is used for the sampling interval of the columns of `Ki` (see “Two-Dimensional Example” on page 4-76). On the first iteration `S` is NaN, so that some initial sampling interval must be determined by `seminfcon`.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `fseminf`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Describes the exit condition:
<code>> 0</code>	The function converged to a solution <code>x</code> .
<code>0</code>	The maximum number of function evaluations or iterations was exceeded.
<code>< 0</code>	The function did not converge to a solution.
<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>algorithm</code>	Algorithm used.
<code>stepsize</code>	Final step size taken.

Options

Optimization options parameters used by `fseminf`. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information:

DerivativeCheck	Compare user-supplied derivatives (gradients) to finite-differencing derivatives.
Diagnostics	Print diagnostic information about the function to be minimized or solved.
DiffMaxChange	Maximum change in variables for finite-difference gradients.
DiffMinChange	Minimum change in variables for finite-difference gradients.
Display	Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.
GradObj	Gradient for the objective function defined by user. See the description of fun above to see how to define the gradient in fun. The gradient <i>must</i> be provided to use the large-scale method. It is optional for the medium-scale method.
MaxFunEvals	Maximum number of function evaluations allowed.
MaxIter	Maximum number of iterations allowed.
TolCon	Termination tolerance on the constraint violation.
TolFun	Termination tolerance on the function value.
TolX	Termination tolerance on x .

Notes

The recommended sampling interval, S , set in `seminfcon`, may be varied by the optimization routine `fseminf` during the computation because values other than the recommended interval may be more appropriate for efficiency or robustness. Also, the finite region w_j , over which $K_j(x, w_j)$ 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 $K_j(x, w_j)$.

Examples

One-Dimensional Example

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. Since the constraints must be in the form $K_f(x, w_j) \leq 0$ 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 an M-file that computes the objective function.

```
function f = myfun(x,s)
% Objective function
f = sum((x-0.5).^2);
```

Second, write an M-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 =
    0.6673
    0.3013
    0.4023

```

The function value and the maximum values of the semi-infinite constraints at the solution x are

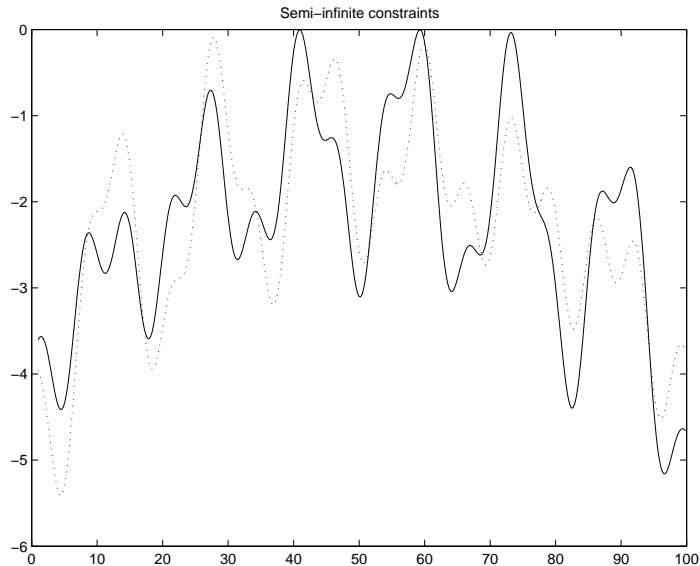
```

fval =
    0.0770

[c,ceq,K1,K2] = mycon(x,NaN); % Use initial sampling interval
max(K1)
ans =
   -0.0017
max(K2)
ans =
   -0.0845

```

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 of 'mycon.m' slows down the computation. Remove this line to improve the speed.

Two-Dimensional Example

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 an M-file that computes the objective function.

```
function f = myfun(x,s)
% Objective function
f = sum((x-0.2).^2);
```

Second, write an M-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
```

fseminf

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 =
    0.2926    0.1874    0.2202
```

and the function value at the solution is


```
fval =
    0.0091
```

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 =
   -0.0027
```

This call to `mycon` produces the following surf plot, which shows the semi-infinite constraint at `x`.



Algorithm

`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 function `fmincon`. 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 more or fewer points need to be taken. The effectiveness of the interpolation is also taken into consideration 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.

See also “SQP Implementation” on page 2-30 in the “Introduction to Algorithms” section for more details on the algorithm used and the types of

fseminf

procedures printed under the Procedures heading for
`options.Display = 'iter'` setting.

Limitations

The function to be minimized, the constraints, and semi-infinite constraints, must be continuous functions of x and w . `fseminf` may only give local solutions.

When the problem is not feasible, `fseminf` attempts to minimize the maximum constraint value.

See Also

`@(function_handle), fmincon, optimset`

Purpose Solve a system of nonlinear equations

$$F(x) = 0$$

for x , where x is a vector and $F(x)$ is a function that returns a vector value.

Syntax

```
x = fsolve(fun,x0)
x = fsolve(fun,x0,options)
x = fsolve(fun,x0,options,P1,P2, ... )
[x,fval] = fsolve(...)
[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.

`x = fsolve(fun,x0)` starts at `x0` and tries to solve the equations described in `fun`.

`x = fsolve(fun,x0,options)` minimizes with the optimization parameters specified in the structure options.

`x = fsolve(fun,x0,options,P1,P2,...)` passes the problem-dependent parameters `P1`, `P2`, etc., directly to the function `fun`. Pass an empty matrix for options to use the default values for options.

`[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`.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `fsolve`. This section provides function-specific details for `fun` and `options`:

`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.

```
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 an inline object.

```
x = fsolve(inline('sin(x.*x)'),x0);
```

If the Jacobian can also be computed *and* `options.Jacobian` is 'on', set by

```
options = optimset('Jacobian','on')
```

then the function `fun` must return, in a second output argument, the Jacobian value `J`, a matrix, at `x`. Note that 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`, 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)`. (Note that the Jacobian `J` is the transpose of the gradient of `F`.)

`options` “Options” on page 4-83 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `fsolve`. This section provides function-specific details for `exitflag` and `output`:

`exitflag` Describes the exit condition:

- > 0 The function converged to a solution x .
- 0 The maximum number of function evaluations or iterations was exceeded.
- < 0 The function did not converge to a solution.

`output` Structure containing information about the optimization. The fields of the structure are:

- `iterations` Number of iterations taken.
- `funcCount` Number of function evaluations.
- `algorithm` Algorithm used.
- `cgiterations` Number of PCG iterations (large-scale algorithm only).
- `stepsize` Final step size taken (medium-scale algorithm only).
- `firstorderopt` Measure of first-order optimality (large-scale algorithm only).
For large scale problems, the first-order optimality is the infinity norm of the gradient $g = J^T F$ (see “Nonlinear Least Squares” on page 3-11).

Options

Optimization options parameters used by `fsolve`. Some parameters apply to all algorithms, some are only relevant when using the large-scale algorithm, and others are only relevant when using the medium-scale algorithm. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information.

We start by describing the `LargeScale` option since it states a *preference* for which algorithm to use. It is only a preference since certain conditions must be

met to use the large-scale algorithm. For `fsolve`, 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` or else the medium-scale algorithm is used:

`LargeScale` Use large-scale algorithm if possible when set to 'on'.
 Use medium-scale algorithm when set to 'off'.

Medium-Scale and Large-Scale Algorithms. These parameters are used by both the medium-scale and large-scale algorithms:

`Diagnostics` Print diagnostic information about the function to be minimized.

`Display` Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.

`Jacobian` If 'on', `fsolve` uses a user-defined Jacobian (defined in `fun`), or Jacobian information (when using `JacobMult`), for the objective function. If 'off', `fsolve` approximates the Jacobian using finite differences.

`MaxFunEvals` Maximum number of function evaluations allowed.

`MaxIter` Maximum number of iterations allowed.

`TolFun` Termination tolerance on the function value.

`TolX` Termination tolerance on `x`.

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

`JacobMult` Function handle for Jacobian multiply function. For large-scale structured problems, this function computes the Jacobian matrix products J^*Y , $J^T Y$, or $J^T (J^*Y)$ without actually forming J . The function is of the form

$$W = \text{jmfun}(Jinfo, Y, flag, p1, p2, \dots)$$

where J_{info} and the additional parameters p_1, p_2, \dots contain the matrices used to compute $J*Y$ (or $J'*Y$, or $J'*(J*Y)$). The first argument J_{info} must be the same as the second argument returned by the objective function fun .

$$[F, J_{\text{info}}] = \text{fun}(x, p_1, p_2, \dots)$$

The parameters p_1, p_2, \dots are the same additional parameters that are passed to `fsolve` (and to `fun`).

$$\text{fsolve}(\text{fun}, \dots, \text{options}, p_1, p_2, \dots)$$

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. `fsolve` uses J_{info} to compute the preconditioner.

Note 'Jacobian' must be set to 'on' for J_{info} to be passed from `fun` to `jmfun`.

See “Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints” on page 1-48 for a similar example.

JacobPattern

Sparsity pattern of the Jacobian for finite-differencing. If it is not convenient to compute the Jacobian matrix J in `fun`, `lsqnonlin` can approximate J via sparse finite-differences provided the structure of J — i.e., locations of the nonzeros — is supplied as the value for `JacobPattern`. In the worst case, if the structure is unknown, you can set `JacobPattern` to be a dense matrix and a full finite-difference approximation is computed in each iteration (this is the default if `JacobPattern` is not set). This can be very expensive for large problems so it is usually worth the effort to determine the sparsity structure.

MaxPCGIter	Maximum number of PCG (preconditioned conjugate gradient) iterations (see the <i>Algorithm</i> section below).
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.
TolPCG	Termination tolerance on the PCG iteration.
TypicalX	Typical x values.

Medium-Scale Algorithm Only. These parameters are used only by the medium-scale algorithm:

DerivativeCheck	Compare user-supplied derivatives (Jacobian) to finite-differencing derivatives.
DiffMaxChange	Maximum change in variables for finite-differencing.
DiffMinChange	Minimum change in variables for finite-differencing.
LevenbergMarquardt	Choose Levenberg-Marquardt over Gauss-Newton algorithm.
LineSearchType	Line search algorithm choice.

Examples

Example 1. This example finds a zero of the system of two equations and two unknowns

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

Thus we want to solve the following system for x

$$\begin{aligned} 2x_1 - x_2 - e^{-x_1} &= 0 \\ -x_1 + 2x_2 - e^{-x_2} &= 0 \end{aligned}$$

starting at $x_0 = [-5 \ -5]$.

First, write an M-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))];
```

Next, call an optimization routine.

```
x0 = [-5; -5];           % Make a starting guess at the solution
options=optimset('Display','iter'); % Option to display output
[x,fval] = fsolve(@myfun,x0,options) % Call optimizer
```

After 28 function evaluations, a zero is found.

Iteration	Func-count	f(x)	Norm of step	First-order optimality	CG-iterations
1	4	47071.2	1	2.29e+004	0
2	7	6527.47	1.45207	3.09e+003	1
3	10	918.372	1.49186	418	1
4	13	127.74	1.55326	57.3	1
5	16	14.9153	1.57591	8.26	1
6	19	0.779051	1.27662	1.14	1
7	22	0.00372453	0.484658	0.0683	1
8	25	9.21617e-008	0.0385552	0.000336	1
9	28	5.66133e-017	0.000193707	8.34e-009	1

Optimization terminated successfully:

Relative function value changing by less than OPTIONS.TolFun

```
x =
    0.5671
    0.5671
```

```
fval =
1.0e-008 *
   -0.5320
   -0.5320
```

Optimization terminated successfully:

Relative function value changing by less than OPTIONS.TolFun

```
x =
    0.5671
    0.5671
fval =
1.0e-08 *
   -0.5320
   -0.5320
```

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 an M-file that computes the equations to be solved.

```
function F = myfun(x)
F = x*x*x-[1,2;3,4];
```

Next, invoke an optimization routine.

```
x0 = ones(2,2); % Make a starting guess at the solution
options = optimset('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-03 *
     0.1541   -0.1163
     0.0109   -0.0243
exitflag =
     1
```

and the residual is close to zero.

```
sum(sum(Fval.*Fval))
ans =
    3.7974e-008
```

Notes

If the system of equations is linear, then `\` (the backslash operator; see `help slash`) should be used for better speed and accuracy. For example, to find the solution to the following linear system of equations.

$$\begin{aligned}3x_1 + 11x_2 - 2x_3 &= 7 \\x_1 + x_2 - 2x_3 &= 4 \\x_1 - x_2 + x_3 &= 19\end{aligned}$$

Then the problem is formulated and solved 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
```

Algorithm

The methods are based on the nonlinear least squares algorithms also used in `lsqnonlin`. The advantage of using a least squares method is that if the system of equations is never zero due to small inaccuracies, or because it just does 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 may converge to a point that is not a solution of the system of equations (see “Limitations” and “Diagnostics” below).

Large-Scale Optimization. By default `fsolve` chooses the large-scale algorithm. The algorithm is a subspace trust region method and is based on the interior-reflective Newton method described in [1],[2]. 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 3-3, and “Preconditioned Conjugate Gradients” on page 3-6 in the “Large-Scale Algorithms” section.

Medium-Scale Optimization. `fsolve` with `options.LargeScale` set to 'off' uses the Gauss-Newton method [3] with line-search. Alternatively, a Levenberg-Marquardt method [4], [5], [6] with line-search may be selected. The choice of algorithm is made by setting `options.LevenbergMarquardt`. Setting `options.LevenbergMarquardt` to 'on' (and `options.LargeScale` to 'off') selects the Levenberg-Marquardt method.

The default line search algorithm, i.e., `options.LineSearchType` set to 'quadcubic', is a safeguarded mixed quadratic and cubic polynomial

interpolation and extrapolation method. A safeguarded cubic polynomial method can be selected by setting `options.LineSearchType` to `'cubicpoly'`. This method generally requires fewer function evaluations but more gradient evaluations. Thus, if gradients are being supplied and can be calculated inexpensively, the cubic polynomial line search method is preferable. The algorithms used are described fully in “Standard Algorithms” on page 2-1.

Diagnostics

`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.

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.

Large-Scale Optimization. Currently, if the analytical Jacobian is provided in `fun`, the `options` parameter `DerivativeCheck` cannot be used with the large-scale method to compare the analytic Jacobian to the finite-difference Jacobian. Instead, use the medium-scale method to check the derivative with `options` parameter `MaxIter` set to 0 iterations. Then run the problem again with the large-scale method. See Table 1-4, Large-Scale Problem Coverage and Requirements for more information on what problem formulations are covered and what information must be provided.

The preconditioner computation used in the preconditioned conjugate gradient part of the large-scale 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 $J^T J$, may lead to a costly solution process for large problems.

See Also

@(function_handle), \, inline, lsqcurvefit, lsqnonlin, optimset

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] 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.

fzero

Purpose Zero of a continuous function of one variable

Syntax

```
x = fzero(fun,x0)
x = fzero(fun,x0,options)
x = fzero(fun,x0,options,P1,P2,...)
[x,fval] = fzero(...)
[x,fval,exitflag] = fzero(...)
[x,fval,exitflag,output] = fzero(...)
```

Description `x = fzero(fun,x0)` tries to find a zero of `fun` near `x0`, if `x0` is a scalar. The value `x` returned by `fzero` is near a point where `fun` changes sign, or NaN if the search fails. In this case, the search terminates when the search interval is expanded until an Inf, NaN, or complex value is found.

If `x0` is a vector of length two, `fzero` assumes `x0` is an interval where the sign of `fun(x0(1))` differs from the sign of `fun(x0(2))`. An error occurs if this is not true. Calling `fzero` with such an interval guarantees `fzero` returns a value near a point where `fun` changes sign.

Note Calling `fzero` with an interval (`x0` with two elements) is often faster than calling it with a scalar `x0`.

`x = fzero(fun,x0,options)` minimizes with the optimization parameters specified in the structure options.

`x = fzero(fun,x0,options,P1,P2,...)` provides for additional arguments, `P1`, `P2`, etc., which are passed to the objective function, `fun`. Use `options = []` as a placeholder if no options are set.

`[x,fval] = fzero(...)` returns the value of the objective function `fun` at the solution `x`.

`[x,fval,exitflag] = fzero(...)` returns a value `exitflag` that describes the exit condition.

`[x,fval,exitflag,output] = fzero(...)` returns a structure output that contains information about the optimization.

Note For the purposes of this command, zeros are considered to be points where the function actually crosses, not just touches, the x -axis.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `fzero`. This section provides function-specific details for `fun` and options:

fun The function whose zero is to be computed. `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.

```
x = fzero(@myfun,x0)
```

where `myfun` is a MATLAB function such as

```
function f = myfun(x)
f = ...           % Compute function value at x
```

`fun` can also be an inline object.

```
x = fzero(inline('sin(x*x)'),x0);
```

options Optimization parameter options. You can set or change the values of these parameters using the `optimset` function. `fzero` uses these options structure fields:

Display Level of display. '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.

TolX Termination tolerance on `x`.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `fzero`. This section provides function-specific details for `exitflag` and output:

exitflag Describes the exit condition:

> 0 Indicates that `fzero` found a zero `x`.

fzero

< 0 No interval was found with a sign change, or a NaN or Inf function value was encountered during the search for an interval containing a sign change, or a complex function value was encountered during the search for an interval containing a sign change.

output Structure containing information about the optimization. The fields of the structure are:

 iterations Number of iterations taken (for fzero, this is the same as the number of function evaluations).

 funcCount Number of function evaluations.

 algorithm Algorithm used.

Examples

Calculate π by finding the zero of the sine function near 3.

```
x = fzero(@sin,3)
x =
    3.1416
```

To find the zero of cosine between 1 and 2,

```
x = fzero(@cos,[1 2])
x =
    1.5708
```

Note that $\cos(1)$ and $\cos(2)$ differ in sign.

To find a zero of the function

$$f(x) = x^3 - 2x - 5$$

write an M-file called `f.m`.

```
function y = f(x)
y = x.^3-2*x-5;
```

To find the zero near 2

```
z = fzero(@f,2)
z =
    2.0946
```

Since this function is a polynomial, the statement `roots([1 0 -2 -5])` finds the same real zero, and a complex conjugate pair of zeros.

```
2.0946
-1.0473 + 1.1359i
-1.0473 - 1.1359i
```

Algorithm

The `fzero` command is an M-file. The algorithm, which was originated 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 the `fzero` M-file is based, is in [2].

Limitations

The `fzero` command finds a point where the function changes sign. If the function is *continuous*, this is also a point where the function has a value near zero. If the function is not continuous, `fzero` may return values that are discontinuous points instead of zeros. For example, `fzero(@tan, 1)` returns 1.5708, a discontinuous point in `tan`.

Furthermore, the `fzero` command defines a *zero* as a point where the function crosses the x -axis. Points where the function touches, but does not cross, the x -axis are not valid zeros. For example, $y = x.^2$ is a parabola that touches the x -axis at 0. Since the function never crosses the x -axis, however, no zero is found. For functions with no valid zeros, `fzero` executes until `Inf`, `NaN`, or a complex value is detected.

See Also

`@(function_handle)`, `\`, `fminbnd`, `fsolve`, `inline`, `optimset`, `roots`

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.

linprog

Purpose Solve a linear programming problem

$$\begin{aligned} \min_x f^T x \quad \text{such that} \quad & A \cdot x \leq b \\ & A_{eq} \cdot x = beq \\ & lb \leq x \leq ub \end{aligned}$$

where f , x , b , beq , lb , and ub are vectors and A and A_{eq} are matrices.

Syntax

```
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,fval] = linprog(...)
[x,fval,exitflag] = linprog(...)
[x,fval,exitflag,output] = linprog(...)
[x,fval,exitflag,output,lambda] = linprog(...)
```

Description

`linprog` solves linear programming problems.

`x = linprog(f,A,b)` solves $\min f^T x$ such that $A^*x \leq b$.

`x = linprog(f,A,b,Aeq,beq)` solves the problem above while additionally satisfying the equality constraints $A_{eq}^*x = beq$. Set $A=[]$ and $b=[]$ if no inequalities exist.

`x = 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 $A_{eq}=[]$ and $beq=[]$ if no equalities exist.

`x = linprog(f,A,b,Aeq,beq,lb,ub,x0)` sets the starting point to x_0 . This option is only available with the medium-scale algorithm (options.LargeScale is 'off'). The default large-scale algorithm ignores any starting point.

`x = linprog(f,A,b,Aeq,beq,lb,ub,x0,options)` minimizes with the optimization parameters specified in the structure options.

`[x,fval] = linprog(...)` returns the value of the objective function `fun` at the solution `x`: `fval = f'*x`.

`[x,lambda,exitflag] = linprog(...)` returns a value `exitflag` that describes the exit condition.

`[x,lambda,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`.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `linprog`. “Options” on page 4-98 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `linprog`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

`exitflag` Describes the exit condition:

- > 0 The function converged to a solution `x`.
- 0 The maximum number of function evaluations or iterations was exceeded.
- < 0 The function did not converge to a solution.

`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 taken.

`algorithm` Algorithm used.
`cgiterations` The number of PCG iterations (large-scale algorithm only).

Options

Optimization options parameters used by `linprog`. Some parameters apply to all algorithms, and others are only relevant when using the large-scale algorithm. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information.:

`LargeScale` Use large-scale algorithm when set to 'on'. Use medium-scale algorithm when set to 'off'.

Medium-Scale and Large-Scale Algorithms. These parameters are used by both the medium-scale and large-scale algorithms:

`Diagnostics` Print diagnostic information about the function to be minimized.

`Display` Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output. At this time, the 'iter' level only works with the large-scale algorithm.

`MaxIter` Maximum number of iterations allowed.

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

`ToIFun` Termination tolerance on the function value.

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);
```

Entering `x`, `lambda.ineqlin`, and `lambda.lower` gets

```

x =
    0.0000
   15.0000
    3.0000
lambda.ineqlin =
    0
    1.5000
    0.5000
lambda.lower =
    1.0000
    0
    0

```

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).

Algorithm

Large-Scale Optimization. The large-scale 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 "Large-Scale Linear Programming" on page 3-14.

Medium-Scale Optimization. `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]. It finds an initial feasible solution by first solving another linear programming problem.

Diagnostics

Large-Scale Optimization. The first stage of the algorithm may involve some preprocessing of the constraints (see "Large-Scale Linear Programming" on page 3-14 in the "Large-Scale Algorithms" section). 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 to not 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 algorithm begins until the stopping criteria is met. (See “Large-Scale Linear Programming” on page 3-14 of the “Large-Scale Algorithms” section 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 displays, 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 displays, it is followed by one of the following six messages indicating if it appears that the dual, the primal, or both are 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 $> \sqrt{\text{TolFun}}$. (The dual residual $< 10 * \text{TolFun}$.)

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.

Medium-Scale Optimization. linprog gives a warning when the solution 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

Medium-Scale Optimization. At this time, the only levels of display, using the Display parameter in options, are 'off' and 'final'; iterative output using 'iter' is not available.

See Also quadprog

References

- [1] Dantzig, G.B., A. Orden, and P. Wolfe, "Generalized Simplex Method for Minimizing a Linear from Under Linear Inequality Constraints," *Pacific Journal Math.* Vol. 5, pp. 183-195.
- [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.

lsqcurvefit

Purpose

Solve nonlinear curve-fitting (data-fitting) problems in the least squares sense. That is, given input data $xdata$, and the observed output $ydata$, find coefficients x that “best-fit” the equation $F(x, xdata)$

$$\min_x \frac{1}{2} \|F(x, xdata) - ydata\|_2^2 = \frac{1}{2} \sum_i (F(x, xdata_i) - ydata_i)^2$$

where $xdata$ and $ydata$ are vectors and $F(x, xdata)$ is a vector valued function.

The function `lsqcurvefit` uses the same algorithm as `lsqnonlin`. Its purpose is to provide an interface designed specifically 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(fun,x0,xdata,ydata,lb,ub,options,P1,P2,...)
[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

`lsqcurvefit` solves nonlinear data-fitting problems. `lsqcurvefit` requires a user-defined function to compute the vector-valued function $F(x, xdata)$. The size of the vector returned by the user-defined function must be the same as the size of $ydata$.

`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`.

`x = lsqcurvefit(fun,x0,xdata,ydata,lb,ub)` defines a set of lower and upper bounds on the design variables, x , so that the solution is always in the range $lb \leq x \leq ub$.

`x = lsqcurvefit(fun,x0,xdata,ydata,lb,ub,options)` minimizes with the optimization parameters specified in the structure options. Pass empty matrices for `lb` and `ub` if no bounds exist.

`x = lsqcurvefit(fun,x0,xdata,ydata,lb,ub,options,P1,P2,...)` passes the problem-dependent parameters `P1`, `P2`, etc., directly to the function `fun`. Pass an empty matrix for options to use the default values for options.

`[x,resnorm] = lsqcurvefit(...)` returns the value of the squared 2-norm of the residual at `x`: $\sum\{(\text{fun}(x,\text{xdata})-\text{ydata})^2\}$.

`[x,resnorm,residual] = lsqcurvefit(...)` returns the value of the residual, $\text{fun}(x,\text{xdata})-\text{ydata}$, 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`.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `lsqcurvefit`. This section provides function-specific details for `fun` and options:

fun The function to be fit. **fun** is a function that takes a vector x and returns a vector F , the objective functions evaluated at x . The function **fun** can be specified as a function handle.

```
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
```

fun can also be an inline object.

```
f = inline('x(1)*xdata.^2+x(2)*sin(xdata)',...
           'x','xdata');
x = lsqcurvefit(f,x0,xdata,ydata);
```

If the Jacobian can also be computed *and* options.Jacobian is 'on', set by

```
options = optimset('Jacobian','on')
```

then the function **fun** must return, in a second output argument, the Jacobian value J , a matrix, at x . Note that 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 nargin > 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 x_0 , then the Jacobian J is an m -by- n matrix where $J(i,j)$ is the partial derivative of $F(i)$ with respect to $x(j)$. (Note that the Jacobian J is the transpose of the gradient of F .)

options “Options” on page 4-108 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `lsqcurvefit`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Describes the exit condition:
> 0	The function converged to a solution x .
0	The maximum number of function evaluations or iterations was exceeded.
< 0	The function did not converge to a solution.
<code>lambda</code>	Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields of the structure are:
lower	Lower bounds <code>lb</code>
upper	Upper bounds <code>ub</code>
<code>output</code>	Structure containing information about the optimization. The fields of the structure are:
iterations	Number of iterations taken.
funcCount	Number of function evaluations.
algorithm	Algorithm used.
cgiterations	The number of PCG iterations (large-scale algorithm only).
stepsize	The final step size taken (medium-scale algorithm only).
firstorderopt	Measure of first-order optimality (large-scale algorithm only)
	For large-scale bound constrained problems, the first-order optimality is the infinity norm of $v .* g$, where v is defined as in “Box Constraints” on page 3-8, and g is the gradient $g = J^T F$ (see “Nonlinear Least Squares” on page 3-11).

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 parameters used by `lsqcurvefit`. Some parameters apply to all algorithms, some are only relevant when using the large-scale algorithm, and others are only relevant when using the medium-scale algorithm. You can use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, Optimization Options Parameters, on page 4-11, for detailed information.

We start by describing the `LargeScale` option since it states a *preference* for which algorithm to use. It is only a preference since certain conditions must be met to use the large-scale or medium-scale algorithm. For the large-scale algorithm, the nonlinear system of equations cannot be under-determined; 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 large-scale algorithm handles bound constraints:

`LargeScale` Use large-scale algorithm if possible when set to 'on'.
 Use medium-scale algorithm when set to 'off'.

Medium-Scale and Large-Scale Algorithms. These parameters are used by both the medium-scale and large-scale algorithms:

`Diagnostics` Print diagnostic information about the function to be minimized.

`Display` Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.

`Jacobian` If 'on', `lsqcurvefit` uses a user-defined Jacobian (defined in `fun`), or Jacobian information (when using `JacobMult`), for the objective function. If 'off', `lsqcurvefit` approximates the Jacobian using finite differences.

`MaxFunEvals` Maximum number of function evaluations allowed.

MaxIter	Maximum number of iterations allowed.
TolFun	Termination tolerance on the function value.
TolX	Termination tolerance on x.

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

JacobMult Function handle for Jacobian multiply function. For large-scale structured problems, this function computes the Jacobian matrix products J^*Y , $J' * Y$, or $J' * (J^*Y)$ without actually forming J . The function is of the form

$$W = \text{jmfun}(J\text{info}, Y, \text{flag}, p1, p2, \dots)$$

where $J\text{info}$ and the additional parameters $p1, p2, \dots$ contain the matrices 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 .

$$[F, J\text{info}] = \text{fun}(x, p1, p2, \dots)$$

The parameters $p1, p2, \dots$ are the same additional parameters that are passed to `lsqcurvefit` (and to fun).

$$\text{lsqcurvefit}(\text{fun}, \dots, \text{options}, p1, p2, \dots)$$

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 $\text{flag} == 0$ then $W = J' * (J^*Y)$. If $\text{flag} > 0$ then $W = J^*Y$. If $\text{flag} < 0$ then $W = J' * Y$. In each case, J is not formed explicitly. `lsqcurvefit` uses $J\text{info}$ to compute the preconditioner.

Note 'Jacobian' must be set to 'on' for Jinfo to be passed from fun to jmfun.

See “Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints” on page 1-48 for a similar example.

JacobPattern	Sparsity pattern of the Jacobian for finite-differencing. If it is not convenient to compute the Jacobian matrix J in fun, lsqcurvefit can approximate J via sparse finite-differences provided the structure of J , i.e., locations of the nonzeros, is supplied as the value for JacobPattern. In the worst case, if the structure is unknown, you can set JacobPattern to be a dense matrix and a full finite-difference approximation is computed in each iteration (this is the default if JacobPattern is not set). This can be very expensive for large problems so it is usually worth the effort to determine the sparsity structure.
MaxPCGIter	Maximum number of PCG (preconditioned conjugate gradient) iterations (see the “Algorithm” section below).
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.
TolPCG	Termination tolerance on the PCG iteration.
TypicalX	Typical x values.

Medium-Scale Algorithm Only. These parameters are used only by the medium-scale algorithm:

DerivativeCheck	Compare user-supplied derivatives (Jacobian) to finite-differencing derivatives.
DiffMaxChange	Maximum change in variables for finite-differencing.
DiffMinChange	Minimum change in variables for finite-differencing.
LevenbergMarquardt	Choose Levenberg-Marquardt over Gauss-Newton algorithm.
LineSearchType	Line search algorithm choice.

Examples

Vectors of data *xdata* and *ydata* are of length *n*. You want to find coefficients *x* to find the best fit to the equation

$$ydata(i) = x(1) \cdot xdata(i)^2 + x(2) \cdot \sin(xdata(i)) + x(3) \cdot xdata(i)^3$$

that is, you want to minimize

$$\min_x \frac{1}{2} \sum_{i=1}^n (F(x, xdata_i) - ydata_i)^2$$

where $F(x, xdata) = x(1) \cdot xdata.^2 + x(2) \cdot \sin(xdata) + x(3) \cdot xdata.^3$, starting at the point $x_0 = [0.3, 0.4, 0.1]$.

First, write an M-file to return the value of F (F has n components).

```
function F = myfun(x,xdata)
F = x(1)*xdata.^2 + x(2)*sin(xdata) + x(3)*xdata.^3;
```

Next, invoke an optimization routine:

```
% Assume you determined xdata and ydata experimentally
xdata = [3.6 7.7 9.3 4.1 8.6 2.8 1.3 7.9 10.0 5.4];
ydata = [16.5 150.6 263.1 24.7 208.5 9.9 2.7 163.9 325.0 54.3];
x0 = [10, 10, 10] % Starting guess
[x, resnorm] = lsqcurvefit(@myfun,x0,xdata,ydata)
```

Note that 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 33 function evaluations, this example gives the solution

```
x =  
0.2269    0.3385    0.3021  
% residual or sum of squares  
resnorm =  
    6.2950
```

The residual is not zero because in this case there was some noise (experimental error) in the data.

Algorithm

Large-Scale Optimization. By default `lsqcurvefit` chooses the large-scale algorithm. This algorithm is a subspace trust region method and is based on the interior-reflective Newton method described in [1], [2]. 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 3-3, and “Preconditioned Conjugate Gradients” on page 3-6 in the “Large-Scale Algorithms” section.

Medium-Scale Optimization. `lsqcurvefit` with `options.LargeScale` set to 'off' uses the Levenberg-Marquardt method with line-search [4], [5], [6]. Alternatively, a Gauss-Newton method [3] with line-search may be selected. The choice of algorithm is made by setting `options.LevenbergMarquardt`. Setting `options.LevenbergMarquardt` to 'off' (and `options.LargeScale` to 'off') selects the Gauss-Newton method, which is generally faster when the residual $\|F(x)\|_2^2$ is small.

The default line search algorithm, i.e., `options.LineSearchType` set to 'quadcubic', is a safeguarded mixed quadratic and cubic polynomial interpolation and extrapolation method. A safeguarded cubic polynomial method can be selected by setting `options.LineSearchType` to 'cubicpoly'. This method generally requires fewer function evaluations but more gradient evaluations. Thus, if gradients are being supplied and can be calculated inexpensively, the cubic polynomial line search method is preferable. The algorithms used are described fully in “Standard Algorithms” on page 2-1.

Diagnostics

Large-Scale Optimization. The large-scale code does not allow equal upper and lower bounds. For example if `lb(2)==ub(2)` then `lsqlin` 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` may only give local solutions.

`lsqcurvefit` only handles real variables (the user-defined function must only return real values). When `x` has complex variables, the variables must be split into real and imaginary parts.

Large-Scale Optimization. The large-scale method for `lsqcurvefit` does not solve underdetermined systems; it requires that the number of equations, i.e., row dimension of F , be at least as great as the number of variables. In the underdetermined case, the medium-scale algorithm is used instead. See Table 1-4, Large-Scale Problem Coverage and Requirements for more information on what problem formulations are covered and what information must be provided.

The preconditioner computation used in the preconditioned conjugate gradient part of the large-scale 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 $J^T J$, may 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.

Currently, if the analytical Jacobian is provided in `fun`, the options parameter `DerivativeCheck` cannot be used with the large-scale method to compare the analytic Jacobian to the finite-difference Jacobian. Instead, use the

medium-scale method to check the derivatives with options parameter `MaxIter` set to zero iterations. Then run the problem with the large-scale method.

See Also `@(function_handle), \lsqlin, lsqnonlin, lsqnonneg, optimset`

- 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.

Purpose Solve the constrained linear least squares problem

$$\min_x \frac{1}{2} \|Cx - d\|_2^2 \quad \text{such that} \quad \begin{aligned} A \cdot x &\leq b \\ Aeq \cdot x &= beq \\ lb &\leq x \leq ub \end{aligned}$$

where C , A , and Aeq are matrices and d , b , beq , lb , ub , and x are vectors.

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(C,d,A,b,Aeq,beq,lb,ub,x0,options,p1,p2,...)
[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

$x = \text{lsqlin}(C,d,A,b)$ 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 .

$x = \text{lsqlin}(C,d,A,b,Aeq,beq)$ solves the problem above while additionally satisfying the equality constraints $Aeq \cdot x = beq$. Set $A=[]$ and $b=[]$ if no inequalities exist.

$x = \text{lsqlin}(C,d,A,b,Aeq,beq,lb,ub)$ 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.

$x = \text{lsqlin}(C,d,A,b,Aeq,beq,lb,ub,x0)$ sets the starting point to $x0$. Set $lb=[]$ and $b=[]$ if no bounds exist.

$x = \text{lsqlin}(C,d,A,b,Aeq,beq,lb,ub,x0,options)$ minimizes with the optimization parameters specified in the structure `options`.

`x = lsqlin(C,d,A,b,Aeq,beq,lb,ub,x0,options,p1,p2,...)` passes the problem-dependent parameters `p1,p2,...` directly to the Jacobian multiply function if it exists. Specify the Jacobian multiply function using the `JacobMult` options parameter.

`[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`.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `lsqlin`. “Options” on page 4-117 provides the function-specific options parameters details.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `lsqlin`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

`exitflag` Describes the exit condition:

- | | |
|-----|--|
| > 0 | The function converged to a solution <code>x</code> . |
| 0 | The maximum number of function evaluations or iterations was exceeded. |
| < 0 | The function did not converge to a solution. |

`lambda` Structure containing the Lagrange multipliers at the solution `x` (separated by constraint type). The fields are:

<code>lower</code>	Lower bounds <code>lb</code>
--------------------	------------------------------

upper	Upper bounds ub
ineqlin	Linear inequalities
eqlin	Linear equalities
output	Structure containing information about the optimization. The fields are:
iterations	Number of iterations taken
algorithm	Algorithm used
cgiterations	Number of PCG iterations (large-scale algorithm only)
firstorderopt	Measure of first-order optimality (large-scale algorithm only) For large-scale bound constrained problems, the first-order optimality is the infinity norm of $v.*g$, where v is defined as in “Box Constraints” on page 3-8, and g is the gradient $g = C^T Cx + C^T d$ (see “Linear Least Squares” on page 3-13).

Options

Optimization options parameters used by `lsqlin`. You can set or change the values of these parameters using the `optimset` function. Some parameters apply to all algorithms, some are only relevant when using the large-scale algorithm, and others are only relevant when using the medium-scale algorithm.

We start by describing the `LargeScale` option since it states a *preference* for which algorithm to use. It is only a preference since certain conditions must be met to use the large-scale algorithm. For `lsqlin`, when the problem has *only* upper and lower bounds, i.e., no linear inequalities or equalities are specified, the default algorithm is the large-scale method. Otherwise the medium-scale algorithm is used:

<code>LargeScale</code>	Use large-scale algorithm if possible when set to 'on'. Use medium-scale algorithm when set to 'off'.
-------------------------	--

Medium-Scale and Large-Scale Algorithms. These parameters are used by both the medium-scale and large-scale algorithms:

Diagnostics	Print diagnostic information about the function to be minimized.
Display	Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.
MaxIter	Maximum number of iterations allowed.

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

JacobMult Function handle for Jacobian multiply function. For large-scale structured problems, this function computes the Jacobian matrix products J^*Y , J'^*Y , or $J'^*(J^*Y)$ without actually forming J . The function is of the form

$$W = \text{jmfun}(J\text{info}, Y, \text{flag}, p1, p2, \dots)$$

where $J\text{info}$ and the additional parameters $p1, p2, \dots$ contain the matrices used to compute J^*Y (or J'^*Y , or $J'^*(J^*Y)$). $J\text{info}$ is the same as the first argument of `lsqlin` and $p1, p2, \dots$ are the same additional parameters that are passed to `lsqlin`.

$$\text{lsqlin}(J\text{info}, \dots, \text{options}, p1, p2, \dots)$$

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 $\text{flag} == 0$ then $W = J'^*(J^*Y)$. If $\text{flag} > 0$ then $W = J^*Y$. If $\text{flag} < 0$ then $W = J'^*Y$. In each case, J is not formed explicitly. `lsqlin` uses $J\text{info}$ to compute the preconditioner.

See “Quadratic Minimization with a Dense but Structured Hessian” on page 1-54 for a related example.

MaxPCGIter Maximum number of PCG (preconditioned conjugate gradient) iterations (see the *Algorithm* section below).

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.
TolFun	Termination tolerance on the function value.
TolPCG	Termination tolerance on the PCG iteration.
TypicalX	Typical x values.

Examples

Find the least squares solution to the over-determined system $C \cdot x = d$ subject to $A \cdot x \leq b$ 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);
```

Entering x, lambda.ineqlin, lambda.lower, lambda.upper gets

```
x =  
-0.1000  
-0.1000  
0.2152  
0.3502  
lambda.ineqlin =  
0  
0.2392  
0  
lambda.lower =  
0.0409  
0.2784  
0  
0  
lambda.upper =  
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, \ should be used. For example, $x = A \setminus b$.

Since 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.

Large-Scale Optimization. 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, then `lsqlin` 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.

Algorithm

Large-Scale Optimization. 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 the large-scale method. 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 3-3 and “Preconditioned Conjugate Gradients” on page 3-6 in the “Large-Scale Algorithms” section.

Medium-Scale Optimization. `lsqlin` with `options.LargeScale` set to `'off'`, or when linear inequalities or equalities are given, 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 “Quadratic Programming” on page 3-12 in the “Introduction to Algorithms” section.

Diagnostics

Large-Scale Optimization. The large-scale code does not allow equal upper and lower bounds. For example if `lb(2) == ub(2)`, then `lsqlin` gives the 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, the medium-scale algorithm must be used to solve equality constrained problems.

Medium-Scale Optimization. If the matrices `C`, `A`, or `Aeq` are sparse, and the problem formulation is not solvable using the large-scale code, `lsqlin` warns that the matrices are converted to full.

```
Warning: This problem formulation not yet available for sparse
matrices.
```

```
Converting to full to solve.
```

`lsqlin` gives a warning when the solution is infeasible.

```
Warning: The constraints are overly stringent;
```

lsqlin

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` parameter in options, are 'off' and 'final'; iterative output using 'iter' is not available.

See Also

`\`, `lsqnonneg`, `quadprog`

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.

Purpose Solve nonlinear least squares (nonlinear data-fitting) problem

$$\min_x f(x) = f_1(x)^2 + f_2(x)^2 + f_3(x)^2 + \dots + f_m(x)^2 + L$$

where L is a constant.

Syntax

```
x = lsqnonlin(fun,x0)
x = lsqnonlin(fun,x0,lb,ub)
x = lsqnonlin(fun,x0,lb,ub,options)
x = lsqnonlin(fun,x0,eb,ub,options,P1,P2,...)
[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 $f(x)$ (the “sum of squares”), `lsqnonlin` requires the user-defined function to compute the *vector*-valued function

$$F(x) = \begin{bmatrix} f_1(x) \\ f_2(x) \\ f_3(x) \end{bmatrix}$$

Then, in vector terms, this optimization problem may be restated as

$$\min_x \frac{1}{2} \|F(x)\|_2^2 = \frac{1}{2} \sum_i f_i(x)^2$$

where x is a vector and $F(x)$ is a function that returns a vector value.

`x = lsqnonlin(fun,x0)` starts at the point x_0 and finds a minimum to the sum of squares of the functions described in `fun`. `fun` should return a vector of

lsqnonlin

values and not the sum-of-squares of the values. ($\text{fun}(x)$ is summed and squared implicitly in the algorithm.)

$x = \text{lsqnonlin}(\text{fun}, x0, \text{lb}, \text{ub})$ defines a set of lower and upper bounds on the design variables, x , so that the solution is always in the range $\text{lb} \leq x \leq \text{ub}$.

$x = \text{lsqnonlin}(\text{fun}, x0, \text{lb}, \text{ub}, \text{options})$ minimizes with the optimization parameters specified in the structure options. Pass empty matrices for lb and ub if no bounds exist.

$x = \text{lsqnonlin}(\text{fun}, x0, \text{lb}, \text{ub}, \text{options}, P1, P2, \dots)$ passes the problem-dependent parameters $P1$, $P2$, etc., directly to the function fun . Pass an empty matrix for options to use the default values for options.

$[x, \text{resnorm}] = \text{lsqnonlin}(\dots)$ returns the value of the squared 2-norm of the residual at x : $\text{sum}(\text{fun}(x).^2)$.

$[x, \text{resnorm}, \text{residual}] = \text{lsqnonlin}(\dots)$ returns the value of the residual, $\text{fun}(x)$, at the solution x .

$[x, \text{resnorm}, \text{residual}, \text{exitflag}] = \text{lsqnonlin}(\dots)$ returns a value exitflag that describes the exit condition.

$[x, \text{resnorm}, \text{residual}, \text{exitflag}, \text{output}] = \text{lsqnonlin}(\dots)$ returns a structure output that contains information about the optimization.

$[x, \text{resnorm}, \text{residual}, \text{exitflag}, \text{output}, \text{lambda}] = \text{lsqnonlin}(\dots)$ returns a structure lambda whose fields contain the Lagrange multipliers at the solution x .

$[x, \text{resnorm}, \text{residual}, \text{exitflag}, \text{output}, \text{lambda}, \text{jacobian}] = \text{lsqnonlin}(\dots)$ returns the Jacobian of fun at the solution x .

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `lsqnonlin`. This section provides function-specific details for `fun` and options:

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.

```
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 an inline object.

```
x = lsqnonlin(inline('sin(x.*x)'),x0);
```

If the Jacobian can also be computed *and* options.**Jacobian** is 'on', set by

```
options = optimset('Jacobian','on')
```

then the function **fun** must return, in a second output argument, the Jacobian value **J**, a matrix, at **x**. Note that 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**, 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)**. (Note that the Jacobian **J** is the transpose of the gradient of **F**.)

options “Options” on page 4-127 provides the function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `lsqnonlin`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Describes the exit condition:
> 0	The function converged to a solution x .
0	The maximum number of function evaluations or iterations was exceeded.
< 0	The function did not converge to a solution.
<code>lambda</code>	Structure containing the Lagrange multipliers at the solution x (separated by constraint type). The fields are:
lower	Lower bounds <code>lb</code>
upper	Upper bounds <code>ub</code>
<code>output</code>	Structure containing information about the optimization. The fields are:
iterations	Number of iterations taken
funcCount	The number of function evaluations
algorithm	Algorithm used
cgiterations	Number of PCG iterations (large-scale algorithm only)
stepsize	The final step size taken (medium-scale algorithm only)
firstorderopt	Measure of first-order optimality (large-scale algorithm only)
	For large-scale bound constrained problems, the first-order optimality is the infinity norm of $v .* g$, where v is defined as in “Box Constraints” on page 3-8, and g is the gradient $g = J^T F$ (see “Nonlinear Least Squares” on page 3-11).

Note The sum of squares should not be formed explicitly. Instead, your function should return a vector of function values. See the example below.

Options

Optimization parameter options. You can set or change the values of these parameters using the `optimset` function. Some parameters apply to all algorithms, some are only relevant when using the large-scale algorithm, and others are only relevant when using the medium-scale algorithm.

We start by describing the `LargeScale` option since it states a *preference* for which algorithm to use. It is only a preference because certain conditions must be met to use the large-scale or medium-scale algorithm. For the large-scale algorithm, the nonlinear system of equations cannot be under-determined; 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 large-scale algorithm handles bound constraints:

`LargeScale` Use large-scale algorithm if possible when set to 'on'.
 Use medium-scale algorithm when set to 'off'.

Medium-Scale and Large-Scale Algorithms. These parameters are used by both the medium-scale and large-scale algorithms:

`Diagnostics` Print diagnostic information about the function to be minimized.

`Display` Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.

`Jacobian` If 'on', `lsqnonlin` uses a user-defined Jacobian (defined in `fun`), or Jacobian information (when using `JacobMult`), for the objective function. If 'off', `lsqnonlin` approximates the Jacobian using finite differences.

`MaxFunEvals` Maximum number of function evaluations allowed.

`MaxIter` Maximum number of iterations allowed.

ToIFun Termination tolerance on the function value.
ToIX Termination tolerance on x .

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

JacobMult Function handle for Jacobian multiply function. For large-scale structured problems, this function computes the Jacobian matrix products J^*Y , $J^'*Y$, or $J^*(J^*Y)$ without actually forming J . The function is of the form

$$W = \text{jmfun}(\text{Jinfo}, Y, \text{flag}, p1, p2, \dots)$$

where Jinfo and the additional parameters $p1, p2, \dots$ contain the matrices 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 .

$$[F, \text{Jinfo}] = \text{fun}(x, p1, p2, \dots)$$

The parameters $p1, p2, \dots$ are the same additional parameters that are passed to `lsqnonlin` (and to fun).

$$\text{lsqnonlin}(\text{fun}, \dots, \text{options}, p1, p2, \dots)$$

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 $\text{flag} == 0$ then $W = J^*(J^*Y)$. If $\text{flag} > 0$ then $W = J^*Y$. If $\text{flag} < 0$ then $W = J^'*Y$. In each case, J is not formed explicitly. `lsqnonlin` uses Jinfo to compute the preconditioner.

Note 'Jacobian' must be set to 'on' for Jinfo to be passed from fun to `jmfun`.

	See “Nonlinear Minimization with a Dense but Structured Hessian and Equality Constraints” on page 1-48 for a similar example.
JacobPattern	Sparsity pattern of the Jacobian for finite-differencing. If it is not convenient to compute the Jacobian matrix J in <code>fun</code> , <code>lsqnonlin</code> can approximate J via sparse finite-differences provided the structure of J , i.e., locations of the nonzeros, is supplied as the value for <code>JacobPattern</code> . In the worst case, if the structure is unknown, you can set <code>JacobPattern</code> to be a dense matrix and a full finite-difference approximation is computed in each iteration (this is the default if <code>JacobPattern</code> is not set). This can be very expensive for large problems so it is usually worth the effort to determine the sparsity structure.
MaxPCGIter	Maximum number of PCG (preconditioned conjugate gradient) iterations (see the <i>Algorithm</i> section below).
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.
TolPCG	Termination tolerance on the PCG iteration.
TypicalX	Typical x values.
Medium-Scale Algorithm Only. These parameters are used only by the medium-scale algorithm:	
DerivativeCheck	Compare user-supplied derivatives (Jacobian) to finite-differencing derivatives.
DiffMaxChange	Maximum change in variables for finite-differencing.
DiffMinChange	Minimum change in variables for finite-differencing.
LevenbergMarquardt	Choose Levenberg-Marquardt over Gauss-Newton algorithm.
LineSearchType	Line search algorithm choice.

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 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 k components).

First, write an M-file to compute the k -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 =
    0.2578    0.2578
resnorm % Residual or sum of squares
resnorm =
    124.3622
```

Algorithm

Large-Scale Optimization. By default `lsqnonlin` chooses the large-scale algorithm. This algorithm is a subspace trust region method and is based on the interior-reflective Newton method described in [1], [2]. 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 3-3 and “Preconditioned Conjugate Gradients” on page 3-6 in the “Large-Scale Algorithms” section.

Medium-Scale Optimization. `lsqnonlin` with `options.LargeScale` set to 'off' uses the Levenberg-Marquardt method with line-search [4], [5], [6]. Alternatively, a Gauss-Newton method [3] with line-search may be selected. The choice of algorithm is made by setting `options.LevenbergMarquardt`. Setting `options.LevenbergMarquardt` to 'off' (and `options.LargeScale` to 'off') selects the Gauss-Newton method, which is generally faster when the residual $\|F(x)\|_2^2$ is small.

The default line search algorithm, i.e., `options.LineSearchType` set to 'quadcubic', is a safeguarded mixed quadratic and cubic polynomial interpolation and extrapolation method. A safeguarded cubic polynomial method can be selected by setting `options.LineSearchType` to 'cubicpoly'. This method generally requires fewer function evaluations but more gradient evaluations. Thus, if gradients are being supplied and can be calculated inexpensively, the cubic polynomial line search method is preferable. The algorithms used are described fully in “Standard Algorithms” on page 2-1.

Diagnostics

Large-Scale Optimization. The large-scale code does not allow equal upper and lower bounds. For example if `lb(2)==ub(2)` then `lsqnonlin` 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` may only give local solutions.

`lsqnonlin` only handles real variables. When x has complex variables, the variables must be split into real and imaginary parts.

Large-Scale Optimization. The large-scale method for `lsqnonlin` does not solve under-determined systems; it requires that the number of equations (i.e., the number of elements of F) be at least as great as the number of variables. In the under-determined case, the medium-scale algorithm is used instead. (If bound constraints exist, a warning is issued and the problem is solved with the bounds ignored.) See Table 1-4, Large-Scale Problem Coverage and

Requirements, for more information on what problem formulations are covered and what information must be provided.

The preconditioner computation used in the preconditioned conjugate gradient part of the large-scale 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 $J^T J$, may lead to a costly solution process for large problems.

If components of x have no upper (or lower) bounds, then `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.

Currently, if the analytical Jacobian is provided in `fun`, the `options` parameter `DerivativeCheck` cannot be used with the large-scale method to compare the analytic Jacobian to the finite-difference Jacobian. Instead, use the medium-scale method to check the derivatives with `options` parameter `MaxIter` set to 0 iterations. Then run the problem with the large-scale method.

Medium-Scale Optimization. The medium-scale algorithm does not handle bound constraints.

Since the large-scale algorithm does not handle under-determined systems and the medium-scale does not handle bound constraints, problems with both these characteristics cannot be solved by `lsqnonlin`.

See Also

`@(function_handle), lsqcurvefit, lsqlin, optimset`

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.

lsqnonneg

Purpose Solves the nonnegative least squares problem

$$\min_x \frac{1}{2} \|Cx - d\|_2^2 \quad \text{such that} \quad x \geq 0$$

where the matrix C and the vector d are the coefficients of the objective function. The vector, x , of independent variables is restricted to be nonnegative.

Syntax

```
x = lsqnonneg(C,d)
x = lsqnonneg(C,d,x0)
x = lsqnonneg(C,d,x0,options)
[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 $\text{norm}(C*x-d)$ subject to $x \geq 0$. C and d must be real.

`x = lsqnonneg(C,d,x0)` uses $x0$ as the starting point if all $x0 \geq 0$; otherwise, the default is used. The default start point is the origin (the default is also used when $x0=[]$ or when only two input arguments are provided).

`x = lsqnonneg(C,d,x0,options)` minimizes with the optimization parameters specified in the structure `options`.

`[x,resnorm] = lsqnonneg(...)` returns the value of the squared 2-norm of the residual, $\text{norm}(C*x-d)^2$.

`[x,resnorm,residual] = lsqnonneg(...)` returns the residual, $C*x-d$.

`[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`.

Arguments

Input Arguments. Table 4-1, *Input Arguments*, contains general descriptions of arguments passed in to `lsqnonneg`. This section provides function-specific details for options:

- `options` Use `optimset` to set or change the values of these fields in the parameters structure, `options`. See Table 4-3, *Optimization Options Parameters*, on page 4-11, for detailed information.
- `Display` Level of display. 'off' 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`.

Output Arguments. Table 4-2, *Output Arguments*, contains general descriptions of arguments returned by `lsqnonneg`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

- `exitflag` Describes the exit condition:
- > 0 The function converged to a solution `x`.
 - 0 The maximum number of function evaluations or iterations was exceeded. Increasing the tolerance `TolX` may lead to a solution.
- `lambda` Vector containing the Lagrange multipliers: `lambda(i) <= 0` when `x(i)` is (approximately) 0, and `lambda(i)` is (approximately) 0 when `x(i) > 0`.
- `output` Structure containing information about the optimization. The fields are:
- `iterations` Number of iterations taken
 - `algorithm` Algorithm used

Examples

Compare the unconstrained least squares solution to the `lsqnonneg` solution for a 4-by-2 problem.

lsqnonneg

```
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)] =
    -2.5627    0
     3.1108    0.6929

[norm(C*(C\d)-d), norm(C*lsqnonneg(C,d)-d)] =
    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.

Algorithm

`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 over-determined)

```
[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, `lsqlin` may be faster than `lsqnonneg`, otherwise `lsqnonneg` is generally more efficient.

See Also `\`, `lsqlin`, `optimset`

References [1] Lawson, C.L. and R.J. Hanson, *Solving Least Squares Problems*, Prentice-Hall, Chapter 23, p. 161, 1974.

optimget

Purpose Get optimization options parameter values.

Syntax

```
val = optimget(options,'param')  
val = optimget(options,'param',default)
```

Description `val = optimget(options,'param')` returns the value of the specified parameter in the optimization options structure `options`. You need to type only enough leading characters to define the parameter name uniquely. Case is ignored for parameter names.

`val = optimget(options,'param',default)` returns `default` if the specified parameter 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` optimization options parameter in the structure called `my_options`.

```
val = optimget(my_options,'Display')
```

This statement returns the value of the `Display` optimization options parameter in the structure called `my_options` (as in the previous example) except that if the `Display` parameter is not defined, it returns the value `'final'`.

```
optnew = optimget(my_options,'Display','final');
```

See Also `optimset`

Purpose	Create or edit optimization options parameter structure.
Syntax	<pre>options = optimset('param1',value1,'param2',value2,...) optimset options = optimset options = optimset(optimfun) options = optimset(olddopts,'param1',value1,...) options = optimset(olddopts,newopts)</pre>
Description	<p><code>options = optimset('param1',value1,'param2',value2,...)</code> creates an optimization options parameter structure called <code>options</code>, in which the specified parameters (<code>param</code>) have specified values. Any unspecified parameters are set to <code>[]</code> (parameters with value <code>[]</code> indicate to use the default value for that parameter when <code>options</code> is passed to the optimization function). It is sufficient to type only enough leading characters to define the parameter name uniquely. Case is ignored for parameter names.</p> <p><code>optimset</code> with no input or output arguments displays a complete list of parameters with their valid values.</p> <p><code>options = optimset</code> (with no input arguments) creates an options structure <code>options</code> where all fields are set to <code>[]</code>.</p> <p><code>options = optimset(optimfun)</code> creates an options structure <code>options</code> with all parameter names and default values relevant to the optimization function <code>optimfun</code>.</p> <p><code>options = optimset(olddopts,'param1',value1,...)</code> creates a copy of <code>olddopts</code>, modifying the specified parameters with the specified values.</p> <p><code>options = optimset(olddopts,newopts)</code> combines an existing options structure <code>olddopts</code> with a new options structure <code>newopts</code>. Any parameters in <code>newopts</code> with nonempty values overwrite the corresponding old parameters in <code>olddopts</code>.</p>
Parameters	For more information about individual parameters, see the reference pages for the optimization functions that use these parameters, or Table 4-3, Optimization Options Parameters.

In the lists below, values in {} denote the default value; some parameters have different defaults for different optimization functions and so no values are shown in {}.

You can also view the optimization parameters and defaults by typing `optimset` at the command line.

Optimization parameters used by both large-scale and medium-scale algorithms:

Diagnostics	'on' {'off'}
Display	'off' 'iter' 'final' 'notify'
GradObj	'on' {'off'}
Jacobian	'on' {'off'}
LargeScale	'on' {'off'}
MaxFunEvals	positive integer
MaxIter	positive integer
TolCon	positive scalar
TolFun	positive scalar
TolX	positive scalar

Optimization parameters used by large-scale algorithms only:

Hessian	'on' {'off'}
HessMult	function {[]}
HessPattern	sparse matrix {sparse matrix of all ones}
JacobMult	function {[]}
JacobPattern	sparse matrix {sparse matrix of all ones}
MaxPCGIter	positive integer {the greater of 1 and floor(n/2)} where n is the number of elements in x0, the starting point
PrecondBandWidth	positive integer {0} Inf

TolPCG positive scalar | {0.1}
 TypicalX vector of all ones

Optimization parameters used by medium-scale algorithms only:

DerivativeCheck 'on' | {'off'}
 DiffMaxChange positive scalar | {1e-1}
 DiffMinChange positive scalar | {1e-8}
 GoalsExactAchieve positive scalar integer | {0}
 GradConstr 'on' | {'off'}
 HessUpdate {'bfgs'} | 'dfp' | 'gillmurray' | 'steepdesc'
 LevenbergMarquardt 'on' | {'off'}
 LineSearchType 'cubicpoly' | {'quadcubic'}
 MeritFunction 'singleobj' | {'multiobj'}
 MinAbsMax positive scalar integer | {0}

Examples

This statement creates an optimization options structure called `options` in which the `Display` parameter is set to `'iter'` and the `TolFun` parameter 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` parameter and storing new values in `optnew`.

```
optnew = optimset(options,'TolX',1e-4);
```

This statement returns an optimization options structure `options` that contains all the parameter 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](#)

Purpose Solve the quadratic programming problem

$$\min_x \frac{1}{2}x^T Hx + f^T x \quad \text{such that} \quad \begin{aligned} A \cdot x &\leq b \\ Aeq \cdot x &= beq \\ lb &\leq x \leq ub \end{aligned}$$

where H , A , and Aeq are matrices, and f , b , beq , lb , ub , and x are vectors.

Syntax

```
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(H,f,A,b,Aeq,beq,lb,ub,x0,options,p1,p2,...)
[x,fval] = quadprog(...)
[x,fval,exitflag] = quadprog(...)
[x,fval,exitflag,output] = quadprog(...)
[x,fval,exitflag,output,lambda] = quadprog(...)
```

Description

$x = \text{quadprog}(H, f, A, b)$ returns a vector x that minimizes $1/2*x'*H*x + f'*x$ subject to $A*x \leq b$.

$x = \text{quadprog}(H, f, A, b, Aeq, beq)$ solves the problem above while additionally satisfying the equality constraints $Aeq*x = beq$.

$x = \text{quadprog}(H, f, A, b, Aeq, beq, lb, ub)$ defines a set of lower and upper bounds on the design variables, x , so that the solution is in the range $lb \leq x \leq ub$.

$x = \text{quadprog}(H, f, A, b, Aeq, beq, lb, ub, x0)$ sets the starting point to $x0$.

$x = \text{quadprog}(H, f, A, b, Aeq, beq, lb, ub, x0, options)$ minimizes with the optimization parameters specified in the structure options.

$x = \text{quadprog}(H, f, A, b, Aeq, beq, lb, ub, x0, options, p1, p2, \dots)$ passes parameters $p1, p2, \dots$ to the Hessian multiply function, if it exists, specified using the `HessMult` parameter in the options structure.

quadprog

`[x,fval] = quadprog(...)` returns the value of the objective function at `x`:
 $fval = 0.5*x'*H*x + f'*x$.

`[x,fval,exitflag] = quadprog(...)` returns a value `exitflag` that describes the exit condition of `quadprog`.

`[x,fval,exitflag,output] = quadprog(...)` returns a structure `output` that contains information about the optimization.

`[x,fval,exitflag,output,lambda] = quadprog(...)` returns a structure `lambda` whose fields contain the Lagrange multipliers at the solution `x`.

Arguments

Input Arguments. Table 4-1, Input Arguments, contains general descriptions of arguments passed in to `quadprog`. “Options” on page 4-145 provides function-specific details for the options parameters.

Output Arguments. Table 4-2, Output Arguments, contains general descriptions of arguments returned by `quadprog`. This section provides function-specific details for `exitflag`, `lambda`, and `output`:

<code>exitflag</code>	Describes the exit condition:
<code>> 0</code>	The function converged to a solution <code>x</code> .
<code>0</code>	The maximum number of function evaluations or iterations was exceeded.
<code>< 0</code>	The function did not converge to a solution.
<code>lambda</code>	Structure containing the Lagrange multipliers at the solution <code>x</code> (separated by constraint type). The fields are:
<code>lower</code>	For the lower bounds <code>lb</code>
<code>upper</code>	For the upper bounds <code>ub</code>
<code>ineqlin</code>	For the linear inequalities
<code>eqlin</code>	For the linear equalities
<code>output</code>	Structure containing information about the optimization. The fields are:
<code>iterations</code>	Number of iterations taken

algorithm	Algorithm used
cgiterations	Number of PCG iterations (large-scale algorithm only)
firstorderopt	Measure of first-order optimality (large-scale algorithm only) For large-scale 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 3-8, and g is the gradient. For large scale problems with linear equalities only, the first-order optimality is the 2-norm of the scaled residual ($z = M \backslash r$) of the Reduced Preconditioned Conjugate Gradient calculation. See “Algorithm” on page 3-6 in “Preconditioned Conjugate Gradients,” and also “Linearly Constrained Problems” on page 3-8.

Options

Optimization parameter options. Use `optimset` to set or change the values of these parameters. Some parameters apply to all algorithms, some are only relevant when using the large-scale algorithm, and others are only relevant when using the medium-scale algorithm.

The parameter to set an algorithm preference:

LargeScale	Use large-scale algorithm if possible when set to 'on'. Use medium-scale algorithm when set to 'off'. 'on' is only a <i>preference</i> . If the problem has <i>only</i> upper and lower bounds, i.e., no linear inequalities or equalities are specified, the default algorithm is the large-scale method. Or, if the problem given to <code>quadprog</code> has <i>only</i> linear equalities, i.e., no upper and lower bounds or linear inequalities are specified, and the number of equalities is no greater than the length of x , the default algorithm is the large-scale method. Otherwise the medium-scale algorithm is used.
------------	--

Medium-Scale and Large-Scale Algorithms. These parameters are used by both the medium-scale and large-scale algorithms:

Diagnostics	Print diagnostic information about the function to be minimized.
Display	Level of display. 'off' displays no output; 'iter' displays output at each iteration; 'final' (default) displays just the final output.
MaxIter	Maximum number of iterations allowed.

Large-Scale Algorithm Only. These parameters are used only by the large-scale algorithm:

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

$$W = \text{hmfun}(\text{Hinfo}, Y, p1, p2, \dots)$$

where Hinfo and the additional parameters $p1, p2, \dots$ contain the matrices used to compute $H*Y$. Hinfo is the same as the first argument of `quadprog` and $p1, p2, \dots$ are the same additional parameters that are passed to `quadprog`.

```
quadprog(Hinfo, ..., options, ...  
         p1, p2, ...)
```

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. `quadprog` uses Hinfo to compute the preconditioner.

See “Quadratic Minimization with a Dense but Structured Hessian” on page 1-54 for an example.

MaxPCGIter Maximum number of PCG (preconditioned conjugate gradient) iterations (see the “Algorithm” section below).

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.
TolFun	Termination tolerance on the function value. TolFun is used as the exit criteria for problems with simple lower and upper bounds (lb, ub).
TolPCG	Termination tolerance on the PCG iteration. TolPCG is used as the exit criteria for problems with only equality constraints (Aeq, beq).
TolX	Termination tolerance on x .
TypicalX	Typical x values.

Examples

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

$$\begin{aligned} x_1 + x_2 &\leq 2 \\ -x_1 + 2x_2 &\leq 2 \\ 2x_1 + x_2 &\leq 3 \\ 0 &\leq x_1, \quad 0 \leq x_2 \end{aligned}$$

First, note that this function can be written in matrix notation as

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

where

$$H = \begin{bmatrix} 1 & -1 \\ -1 & 2 \end{bmatrix}, \quad f = \begin{bmatrix} -2 \\ -6 \end{bmatrix}, \quad x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Enter these 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)
```

Next, invoke a quadratic programming routine.

```
[x,fval,exitflag,output,lambda] = quadprog(H,f,A,b,[],[],lb)
```

This generates the solution

```
x =
    0.6667
    1.3333
fval =
   -8.2222
exitflag =
     1
output =
    iterations: 3
    algorithm: 'medium-scale: active-set'
    firstorderopt: []
    cgiterations: []
lambda.ineqlin
ans =
    3.1111
    0.4444
     0
lambda.lower
ans =
     0
     0
```

Nonzero elements of the vectors in the fields of `lambda` indicate active constraints at the solution. In this case, the first and second inequality constraints (in `lambda.ineqlin`) are active constraints (i.e., the solution is on their constraint boundaries). For this problem, all the lower bounds are inactive.

Notes

In general quadprog locates a local solution unless the problem is strictly convex.

Better numerical results are likely if you specify equalities explicitly using `Aeq` and `beq`, instead of implicitly using `lb` and `ub`.

If components of x have no upper (or lower) bounds, then quadprog 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.

Large-Scale Optimization. If you do not supply `x0`, or `x0` is not strictly feasible, quadprog chooses a new strictly feasible (centered) starting point.

If an equality constrained problem is posed and quadprog detects negative curvature, then the optimization terminates because the constraints are not restrictive enough. In this case, `exitflag` is returned with the value `-1`, a message is displayed (unless the options `Display` parameter is `'off'`), and the `x` returned is not a solution but a direction of negative curvature with respect to H .

For problems with simple lower and upper bounds (`lb`, `ub`), quadprog exits based on the value of `TolFun`. For problems with only equality constraints (`Aeq`, `beq`), the exit is based on `TolPCG`. Adjust `TolFun` and `TolPCG` to affect your results. `TolX` is used by both type of problems.

Algorithm

Large-Scale Optimization. When the problem given to quadprog has *only* upper and lower bounds, i.e., no linear inequalities or equalities are specified, the default algorithm is the large-scale method. Or, if the problem given to quadprog has *only* linear equalities, i.e., no upper and lower bounds or linear inequalities are specified, the default algorithm is the large-scale method.

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 3-3 and “Preconditioned Conjugate Gradients” on page 3-6 in the “Large-Scale Algorithms” section.

Medium-Scale Optimization. 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. This method is discussed in “Standard Algorithms” on page 2-1.

Diagnostics

Large-Scale Optimization. The large-scale code does not allow equal upper and lower bounds. For example if $lb(2) == ub(2)$ then quadprog gives the error

```
Equal upper and lower bounds not permitted in this large-scale method.
```

```
Use equality constraints and the medium-scale method instead.
```

If you only have equality constraints you can still use the large-scale method. But if you have both equalities and bounds, you must use the medium-scale method.

Medium-Scale Optimization. When the solution is infeasible, quadprog gives this warning

```
Warning: The constraints are overly stringent;  
         there is no feasible solution.
```

In this case, quadprog produces a result that minimizes the worst case constraint violation.

When the equality constraints are inconsistent, quadprog gives this warning

```
Warning: The equality constraints are overly stringent;  
         there is no feasible solution.
```

Unbounded solutions, which can occur when the Hessian H is negative semidefinite, may result in

```
Warning: The solution is unbounded and at infinity;  
         the constraints are not restrictive enough.
```

In this case, quadprog returns a value of x that satisfies the constraints.

Limitations

At this time the only levels of display, using the `Display` parameter in options, are 'off' and 'final'; iterative output using 'iter' is not available.

The solution to indefinite or negative definite problems is often unbounded (in this case, `exitflag` is returned with a negative value to show a minimum was not found); when a finite solution does exist, quadprog may only give local minima since the problem may be nonconvex.

Large-Scale Optimization. The linear equalities cannot be dependent (i.e., A_{eq} must have full row rank). Note that this means that A_{eq} cannot have more rows than columns. If either of these cases occur, the medium-scale algorithm is called instead. See Table 1-4, Large-Scale Problem Coverage and Requirements for more information on what problem formulations are covered and what information must be provided.

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. and W. Murray, and M.H. Wright, *Practical Optimization*, Academic Press, London, UK, 1981.

A

active constraints 4-99, 4-120, 4-148
 active set method 2-31, 4-43, 4-100, 4-121, 4-149
 arguments, additional 1-18
 attainment factor 4-24
 axis crossing. *See* zero of a function

B

banana function 2-5
 BFGS formula 2-7, 4-43, 4-66
 bisection search 4-95
 bound constraints, large-scale 3-8
 box constraints. *See* bound constraints

C

centering parameter 3-15
 CG. *See* conjugate gradients
 code
 updating to Version 2 syntax 1-76
 complementarity conditions 3-15
 complex values 1-75
 complex variables 4-113, 4-131
 conjugate gradients 3-4
 constrained minimization 4-32
 large-scale example 1-43, 1-47
 medium-scale example 1-10
 constraints
 linear 3-8, 4-42, 4-52
 positive 1-17
 continuous function and gradient methods 2-5
 conventions in our documentation (table) xvii
 convex problem 2-26
 cubic interpolation 2-11
 curve-fitting 4-104
 categories 1-6

functions that apply 4-4

D

data-fitting 4-104
 categories 1-6
 functions that apply 4-4
 dense columns, constraint matrix 3-16
 DFP formula 4-66
 direction of negative curvature 3-4
 discontinuities 1-73, 2-5
 discontinuous problems 4-57, 4-66
 discrete variables 1-74
 dual problem 3-14
 duality gap 3-15

E

ε -constraint method 2-40
 equality constraints
 dense columns 1-60
 medium-scale example 1-16
 equality constraints inconsistent warning, quad-
 prog 4-150
 equality constraints, linear
 large-scale 3-8
 equation solving
 categories 1-5
 functions that apply 4-3
 error, Out of memory. 1-39

F

feasibility conditions 3-15
 feasible point, finding 2-34
 fgoalattain **4-16**

- example 1-27
- fixed variables 3-17
- fixed-step ODE solver 1-24
- fminbnd **4-28**
- fmincon **4-32**
 - large-scale example 1-43, 1-47
 - medium-scale example 1-10
- fminimax **4-45**
 - example 1-24
- fminsearch **4-54**
- fminunc **4-58**
 - large-scale example 1-40
 - medium-scale example 1-9
 - warning messages 1-73
- fseminf **4-68**
- fsolve **4-81**
 - large-scale example 1-34
 - medium-scale example 1-18
- function arguments 4-6
- function discontinuities 1-73
- functions
 - grouped by type of optimization 4-3
 - updating to Version 2 1-76
- fzero **4-92**

G

- Gauss-Newton 3-11
- Gauss-Newton method 2-20, 2-21, 2-25, 4-89,
4-112, 4-131
- global minimum 1-72
- global variables 1-18
- goal attainment 2-42, 4-16
 - example 1-27
- goaldemo 4-24
- golden section search 4-31
- gradient checking, analytic 1-15

- gradient examples 1-13
- gradient function 1-7
- gradient methods 2-5

H

- Hessian modified message 2-31
- Hessian modified twice message 2-31
- Hessian sparsity structure 1-42
- Hessian update 2-12, 2-30

I

- inconsistent constraints 4-102
- indefinite problems 4-150
- infeasible message 2-31
- infeasible problems 4-43
- infeasible solution warning
 - linprog 4-102
 - quadprog 4-150
- infinite loop 1-75
- inline objects 1-70
- input arguments 4-6
- integer variables 1-74
- interior-point linear programming 3-14
- introduction to optimization 2-4
- iterative display 1-67

J

- Jacobian examples 1-34
- Jacobian sparsity pattern 1-36

K

- Kuhn-Tucker equations 2-26

L

Lagrange multipliers
 large-scale linear programming 3-17
large-scale functionality coverage 1-31
large-scale methods
 demos 4-4
 examples 1-30
least squares 2-21
 categories 1-6
 functions that apply 4-4
Levenberg-Marquardt method 2-22, 2-25, 4-89,
 4-112, 4-131
line search 4-66, 4-89, 4-90, 4-112, 4-131
line search strategy 1-7
linear constraints 3-8, 4-42, 4-52
linear equations solve 4-88
linear least squares
 constrained 4-115
 large-scale algorithm 3-13
 large-scale example 1-58
 nonnegative 4-134
 unconstrained 4-120
linear programming 4-96
 implementation 2-34
 large-scale algorithm 3-14
 large-scale example 1-59, 1-60
 problem 2-4
linprog **4-96**
 large-scale example 1-59, 1-60
LIPSOL 3-14
lower bounds 1-12
lsqcurvefit **4-104**
lsqlin **4-115**
 large-scale example 1-58
lsqnonlin **4-123**
 convergence 1-75
 large-scale example 1-36

 medium-scale example 1-21
lsqnonneg **4-134**

M

maximization 1-17
medium-scale methods
 demos 4-5
Mehrotra's predictor-corrector algorithm 3-14,
 3-15
merit function 2-35
minimax examples 1-24
minimax problem, solving 4-45
minimization
 categories 1-4
 functions that apply 4-3
multiobjective optimization 2-36, 4-16
 examples 1-18

N

NCD. *See* Nonlinear Control Design
negative curvature direction 3-4, 3-6
negative definite problems 4-150
Nelder and Mead 2-5
Newton direction
 approximate 3-4
Newton's method 2-5
no update message 2-31
nonconvex problems 4-150
noninferior solution 2-38
Nonlinear Control Design (NCD) Blockset 1-24
nonlinear data-fitting 4-123
nonlinear equations, solving 4-81
nonlinear least squares 2-24, 4-104, 4-123
 large-scale algorithm 3-11
 large-scale example 1-36

nonlinear programming 2-4
nonlinear system of equations, large-scale exam-
ple 1-34
normal equations 3-11, 3-13

O

objective function 1-6
objective function, undefined values 1-73
optimality conditions linear programming 3-14
optimget **4-138**
optimization
 functions for each type 4-3
 introduction 2-4
optimization parameters structure 1-63, 4-138,
 4-139
optimset **4-139**
options parameters
 descriptions 4-11
 possible values 4-140
 utility functions 4-4
Out of memory. error 1-39
output arguments 4-6
output display 1-65
output headings 1-67
 large-scale algorithms 1-67
 medium-scale algorithms 1-65

P

PCG. *See* preconditioned conjugate gradients
preconditioned conjugate gradients 3-4, 3-6,
 3-16
 algorithm 3-6
preconditioner 1-35, 3-6
 banded 1-43
predictor-corrector algorithm 3-15

preprocessing
 linear programming 3-14, 3-17
primal problem 3-14
primal-dual algorithm 3-15
primal-dual interior-point 3-14
projection method 2-31, 4-149

Q

quadprog **4-143**
 large-scale example 1-52
quadratic interpolation 2-10
quadratic programming 2-4, 4-43, 4-143
 large-scale algorithm 3-12
 large-scale example 1-52
quasi-Newton method 2-7, 2-12, 4-66

R

reflective line search 3-12
reflective steps 3-9, 3-10
residual 2-19
Rosenbrock's function 2-5

S

sampling interval 4-73
secular equation 3-3
semi-infinite constraints 4-68
Sherman-Morrison formula 3-16
signal processing example 1-27
simple bounds 1-12
simplex search 2-5, 4-56
Simulink, multiobjective example 1-19
singleton rows 3-17
sparsity pattern Jacobian 1-36
sparsity structure, Hessian 1-42

SQP method 2-27, 2-31, 4-43
steepest descent 4-66
stopping criteria, large-scale linear programming
 3-16
structural rank 3-17
subspace
 determination of 3-4
subspace, two-dimensional 3-4
syntax
 updating to Version 2 1-76
system of equations, solving 4-81

T
trust region 3-3
two-dimensional subspace 3-4
typographical conventions xvii

U
unbounded solutions warning
 linprog 4-102
 quadprog 4-150
unconstrained minimization 4-54, 4-58
 large-scale example 1-40
 medium-scale example 1-9
 one dimensional 4-28
unconstrained optimization 2-5
updating code to Version 2 syntax 1-76
upper bounds 1-12

V
variable-step ODE solver 1-24
Version 2
 changes to calling sequence 1-77
 converting Version 1.5 code 1-76

W
warning
 equality constraints inconsistent, quadprog
 4-150
 infeasible solution, linprog 4-102
 infeasible solution, quadprog 4-150
 stuck at minimum, fsolve 4-90
 unbounded solutions, linprog 4-102
 unbounded solutions, quadprog 4-150
warnings displayed 1-74
weighted sum strategy 2-39

Z
zero curvature direction 3-6
zero finding 4-81
zero of a function, finding 4-92

