

Model-Based Calibration Toolbox™

CAGE User's Guide



MATLAB® & SIMULINK®

R2021a



How to Contact MathWorks



Latest news: www.mathworks.com
Sales and services: www.mathworks.com/sales_and_services
User community: www.mathworks.com/matlabcentral
Technical support: www.mathworks.com/support/contact_us



Phone: 508-647-7000



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

Model-Based Calibration Toolbox™ CAGE User's Guide

© COPYRIGHT 2001–2021 by The MathWorks, Inc.

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

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

Trademarks

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

Patents

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

Revision History

December 2001	Online only	New for Version 1.0 (Release 12.1)
August 2002	Online only	Revised for Version 1.1 (Release 13)
May 2003	Online only	Revised for Version 2.0 (Release 13+)
June 2004	Online only	Revised for Version 2.1 (Release 14)
June 2004	Online only	Revised for Version 2.1.1 (Release 14+)
November 2005	Online only	Revised for Version 3.0 (Release 14SP3+)
September 2006	Online only	Revised for Version 3.1 (Release 2006b)
March 2007	Online only	Version 3.2 (Release 2007a)
September 2007	Online only	Revised for Version 3.3 (Release 2007b)
March 2008	Online only	Revised for Version 3.4 (Release 2008a)
October 2008	Online only	Revised for Version 3.4.1 (Release 2008a+)
October 2008	Online only	Revised for Version 3.5 (Release 2008b)
March 2009	Online only	Revised for Version 3.6 (Release 2009a)
September 2009	Online only	Revised for Version 3.7 (Release 2009b)
March 2010	Online only	Revised for Version 4.0 (Release 2010a)
September 2010	Online only	Revised for Version 4.1 (Release 2010b)
April 2011	Online only	Revised for Version 4.2 (Release 2011a)
September 2011	Online only	Revised for Version 4.3 (Release 2011b)
March 2012	Online only	Revised for Version 4.4 (Release 2012a)
September 2012	Online only	Revised for Version 4.5 (Release 2012b)
March 2013	Online only	Revised for Version 4.6 (Release 2013a)
September 2013	Online only	Revised for Version 4.6.1 (Release 2013b)
March 2014	Online only	Revised for Version 4.7 (Release 2014a)
October 2014	Online only	Revised for Version 4.8 (Release 2014b)
March 2015	Online only	Revised for Version 4.8.1 (Release 2015a)
September 2015	Online only	Revised for Version 5.0 (Release 2015b)
March 2016	Online only	Revised for Version 5.1 (Release 2016a)
September 2016	Online only	Revised for Version 5.2 (Release 2016b)
March 2017	Online only	Revised for Version 5.2.1 (Release 2017a)
September 2017	Online only	Revised for Version 5.3 (Release 2017b)
March 2018	Online only	Revised for Version 5.4 (Release 2018a)
September 2018	Online only	Revised for Version 5.5 (Release 2018b)
March 2019	Online only	Revised for Version 5.6 (Release 2019a)
September 2019	Online only	Revised for Version 5.7 (Release 2019b)
March 2020	Online only	Revised for Version 5.8 (Release 2020a)
September 2020	Online only	Revised for Version 5.9 (Release 2020b)
March 2021	Online only	Revised for Version 5.10 (Release 2021a)

Getting Started

1

What Is CAGE?	1-2
Navigate CAGE	1-4
Processes	1-4
Data Objects	1-4

Variables and Models

2

Set Up Calibrations, Resume Work, or Find Calibration Examples	2-2
Import Models and Calibration Items Using CAGE Import Tool	2-4
Set Up Variables and Constants	2-7
Introducing the Variable Dictionary View	2-7
Importing and Exporting a Variable Dictionary	2-9
Using the Variable Menu	2-9
Using Aliases	2-10
Set Up Models	2-11
Introducing the Models View	2-11
Importing Models	2-13
Adding New Function Models	2-13
Renaming and Editing Models	2-14
Composite Models	2-16
What Are Composite Models?	2-16
Importing from the Model Browser	2-16
Combining Existing CAGE Models	2-19
Viewing Composite Model Properties	2-20
Classification Models	2-23
Regression Models	2-25

3

Create Lookup Tables from a Model	3-2
Import Lookup Tables from a Simulink Model	3-5
Edit Lookup Tables	3-7
Add, Duplicate, and Delete Tables	3-8
Fill Tables by Extrapolation	3-9
Optimize Lookup Table Breakpoints	3-11
Lock Lookup Table Values	3-12
Edit the Lookup Table Graph	3-12
Arithmetic Operations on Lookup Table Values	3-12
Fill a Single Lookup Table From a Model	3-14
Compare and Revert Lookup Tables	3-17
Reset to Previous Lookup Table Versions	3-17
Compare Versions	3-18
Calibration Manager	3-20
Set Up Lookup Tables from a Calibration File	3-20
Set Up Lookup Tables Manually	3-22
Copy Lookup Table Data from Other Sources	3-23
Lookup Table Properties	3-24
Table Properties: General Tab	3-24
Table Properties: Table Values Precision Tab	3-24
Table Properties: Inputs Tab	3-28
Lookup Table Normalizers	3-29
Editing Breakpoints	3-30
Input/Output Display	3-31
Normalizer Display	3-32
Breakpoint Spacing Display	3-32
Invert Lookup Tables	3-35
Inverting One-Dimensional Lookup Tables	3-36
Inverting Two-Dimensional Lookup Tables	3-37
Import and Export Calibrations	3-39
Formats	3-39
Import Calibrations	3-39
Export Calibrations	3-39

Feature Calibrations

4

About Feature Calibrations and Strategies	4-2
--	------------

Process for Calibrating and Filling Feature Tables	4-3
Set Up Feature and Strategy	4-5
Add a Feature	4-5
Import a Strategy from Simulink	4-5
Construct a Strategy	4-9
Export Strategies	4-11
Initialize Feature and Calibration Tables	4-12
Initialize a Feature	4-12
Initialize Breakpoints	4-13
Initialize Table Values	4-14
Optimize Normalizer Breakpoints	4-14
Fill Feature Lookup Tables	4-18
Filling and Optimizing Lookup Table Values	4-18
Saving and Reusing Feature Fill Settings	4-21
Optimize Calibrations with Transient Controllers	4-22
Import Simulink Model	4-22
Run Optimization	4-23
Restart Optimization	4-25
Analyze Results	4-25
Export Calibration	4-26
Validate Feature Fill	4-28
Optimizing Normalizer Breakpoints and Table Values	4-29
Filling Methods	4-29
ReduceError	4-29
ShareAveCurv and ShareCurvThenAve	4-30
Optimizing Table Values	4-31

Feature Calibration for an ECU

5

Calibrate a Feature for an Estimator Subsystem in an ECU	5-2
Set Up Feature and Torque Strategy	5-4
Import Two-Stage Models	5-6
Align Model and Feature Inputs	5-8
Fill Feature and Calibration Tables	5-10
Fill Settings for Torque Table, T	5-10
Fill Settings for Spark Table, F_SPK	5-11
Fill Settings for AFR Table, F_A	5-13
Validate Feature and Export Lookup Tables	5-15
Validate Feature	5-15
Export Calibration Tables	5-17

Create a Tradeoff from an Optimization	6-2
Import Models	6-3
Create Lookup Tables and Tradeoff from Model	6-4
Create Optimization	6-7
Run Optimization	6-9
Fill Lookup Tables	6-11
Analyze Results	6-13
Perform a Tradeoff Calibration	6-15
Set Up a Tradeoff Calibration	6-17
Add a Tradeoff	6-17
Add Tables to a Tradeoff	6-17
Display Models in Tradeoff	6-20
Fill Lookup Tables in a Tradeoff Calibration	6-21
Set Values of Variables	6-25
Set Values for Individual Operating Points	6-25
Setting Values for All Operating Points	6-26
Choose a Table Value at a Specific Operating Point	6-27
Find Maximum, Minimum, or Turning Point of Graphs	6-29
Use Zoom Controls on the Graphs	6-29
Configure Views	6-29
Configure Inputs	6-30
Control Table Extrapolation Regions	6-31
Define a Region	6-31
Clear a Region	6-31
Point-by-Point Model Tradeoffs	6-33
Adding a Point-by-Point Model Tradeoff	6-33
Calibrating Using a Point-by-Point Model Tradeoff	6-33

Optimization Setup

Using Optimization in CAGE	7-2
Overview of Optimization in CAGE	7-2
Parallel Computing in Optimization	7-2
Optimization Problems You Can Solve with CAGE	7-3
Create an Optimization	7-6
Setting Up Optimizations	7-6
Creating Optimizations from Models	7-7
Tools for Common Optimization Tasks	7-9
Optimization Wizard	7-10

Set Up Sum Optimizations	7-16
Overview of Setting Up Sum Optimizations	7-16
Example Problem to Demonstrate Controls for Sum Optimizations	7-17
Selecting Scalar Variables	7-18
Algorithm Restrictions	7-18
Using Application Point Sets	7-20
Set Up Multiobjective Optimizations	7-22
Overview of Setting Up Multiobjective Optimizations	7-22
About the gamultiobj Algorithm	7-23
About the NBI (Normal Boundary Intersection) Algorithm	7-23
Set Up Modal Optimizations	7-25
What Is Modal Optimization?	7-25
Workflow for Modal Optimization	7-25
Creating Modal Optimizations	7-26
Adding Extra Objectives to Modal Optimizations	7-26
Set Up MultiStart Optimizations	7-28
What Is MultiStart Optimization?	7-28
Creating a MultiStart Optimization	7-28
Edit Variable Values	7-30
What Are Variable Values?	7-30
Define Variables Manually	7-30
Import from a Data Set	7-31
Import from Output	7-32
Import from Lookup Table Breakpoints	7-34
Import from Lookup Table Values	7-35
Edit Objectives and Constraints	7-37
Edit Objective	7-37
Edit Constraint	7-38
Soft Constraints	7-41
Run Optimizations	7-43
Edit Optimization Parameters	7-44
fmincon Optimization Parameters	7-44
NBI Optimization Parameters	7-45
GA Optimization Parameters	7-46
Pattern Search Optimization Parameters	7-47
Modal Optimization Parameters	7-48
MultiStart Optimization Parameters	7-48
gamultiobj Optimization Parameters	7-49
Scale Optimization	7-49

Optimization Analysis

8

Using Optimization Results	8-2
Choosing Acceptable Solutions	8-2

Create Sum Optimization from Point Optimization Output	8-3
Exporting to a Data Set	8-3
Custom Fill Function Structure	8-5
Filling Tables from Optimization Results	8-7
Lookup Table Filling from Optimization Results Wizard	8-7
Lookup Table Filling When Optimization Operating Point Inputs Differ from Table Inputs	8-11
Filling Tables Via Data Sets	8-13
Viewing Your Optimization Results	8-14
Navigating the Optimization Output View	8-14
Solution Slice: Results Surface and Results Contour Views	8-14
Solution Slice: Optimization Results Table	8-17
Objective Slice Graphs	8-18
Objective Contour Plot	8-19
Constraint Slice Graphs	8-20
Constraint Summary Table	8-22
Analyzing Point Optimization Output	8-27
Process for Analyzing Optimization Results	8-27
Detecting Local Optima	8-29
Investigating Early Termination of Optimization	8-31
Handling Flat Optima	8-35
Tools for Optimizations with Multiple Solutions	8-38
Analyzing Modal, MultiStart, and Multiobjective Optimizations	8-38
Pareto Slice Table View	8-38
Selected Solution Slice	8-39
Exporting Selected Solutions	8-41
Analyzing Modal Optimization Results	8-42
Viewing and Selecting Modal Optimization Results	8-42
Creating Sum Optimizations from Modal Optimizations	8-45
Filling Tables for Operating Modes	8-45
Analyzing MultiStart Optimization Results	8-48
Viewing and Selecting MultiStart Results	8-48
Creating Sum Optimizations from MultiStart Optimizations	8-50
Analyzing Multiobjective Optimization Results	8-51
Pareto Graphs	8-51
Multiobjective Output Messages	8-52
Interpreting Sum Optimization Output	8-54
Objective Contour Plot	8-54
Constraint Summary	8-55
Optimization Results Table	8-56
Operating Point Indices	8-57
Objective Graphs	8-57
Constraint Graphs	8-58
Weighted Objective Pareto Slice	8-60

9

User-Defined Optimizations	9-2
Introducing User-Defined Optimization	9-2
Implementing Your Optimization Algorithm in CAGE	9-2
About the Worked Example Optimization Algorithm	9-4
Checking User-Defined Optimizations into CAGE	9-6
Create a User-Defined Optimization	9-7
Overview	9-7
Using the Worked Example Optimization	9-8
Creating an Optimization from Your Own Algorithm	9-12
Process Overview	9-12
Step 1: Verify the Algorithm	9-13
Step 2: Create a CAGE Optimization Function	9-14
Step 3: Define the Optimization Options	9-15
Step 4: Add the Algorithm to the Optimization Function	9-17
Step 5: Register Your Optimization Function with CAGE	9-19
Step 6: Verify Your New Optimization	9-20
Optimization Function Reference	9-23
Methods of cgoptimoptions	9-23
Methods of cgoptimstore	9-24
Functions	9-26

Data Sets

10

Use Data Sets Views	10-2
Set Up Data Sets	10-4
How to Set Up Data Sets	10-4
Importing Experimental Data from File	10-4
Importing Data from the Model Browser	10-6
Importing Data from a Table in Your Session	10-7
Merging Data Sets	10-7
Specifying the Factors Manually	10-7
Creating a Factor from the Error Between Factors	10-9
View Data in a Table	10-11
Plot Outputs	10-12
Use Color to Display Information	10-14
Link Factors in a Data Set	10-17
Manipulate Models in Data Set View	10-19

Export Data Sets	10-20
Exporting Data to the Model Browser	10-20
Exporting Data to File	10-20

Surface Viewer

11

View a Model or Strategy	11-2
Set Variable Ranges	11-4
Display Point-by-Point Models in the Surface Viewer	11-4
Display the Model or Feature	11-6
Use Display Options	11-6
Surface	11-6
Contour	11-8
Line	11-8
Single Value	11-9
Multiline	11-9
Table	11-9
Make Movies	11-11
Display Errors	11-12
Introducing Error Displays	11-12
Feature Error Data	11-12
Prediction Error Data	11-12
Print and Export the Display	11-13

Getting Started

This section includes the following topics:

- “What Is CAGE?” on page 1-2
- “Navigate CAGE” on page 1-4

What Is CAGE?

Model-Based Calibration Toolbox contains tools for design of experiment, statistical modeling, and calibration of complex systems. The toolbox has two main user interfaces:

- Model Browser for design of experiment and statistical modeling
- CAGE Browser for analytical calibration

CAGE (CALibration GENeration) is an easy-to-use graphical interface for calibrating lookup tables for your electronic control unit (ECU).

CAGE uses models of the engine control subsystems to calibrate lookup tables. With CAGE, you fill and optimize lookup tables in existing ECU software using models from the Model Browser part of the Model-Based Calibration Toolbox product. From these models, CAGE builds steady-state ECU calibrations. Use CAGE to calibrate lookup tables, tradeoff competing objectives, and validate calibrations.

Use CAGE For	Description
Feature calibration	<p>A feature calibration compares a model of an estimated signal with a lookup table (or algebraic collection of tables) that estimates the same signal in the ECU. CAGE finds the optimum calibration for the lookup tables.</p> <p>For example, a typical engine subsystem controls the spark angle to produce the peak torque; that is, the Maximum Brake Torque (MBT) spark. Using the Model Browser, you can build a statistically sound model of MBT spark, over a range of engine speeds and relative air charges, or loads. Use the feature calibration to fill a lookup table by comparing the table to the model.</p>
Tradeoff calibration	<p>A tradeoff calibration fills lookup tables by comparing models of different engine characteristics at key operating points.</p> <p>For example, there are several models of important engine characteristics, such as torque and nitrous oxides (NOX) emissions. Both models depend on the spark angle. At a particular operating point, a slight reduction of torque can result in a dramatic reduction of NOX emissions. Thus, the calibrator uses the value of the spark angle that gives this reduction in NOX emissions instead of the spark angle that generates maximum torque.</p>
Optimization	<p>CAGE can optimize calibrations with respect to models, including single- and multi-objective optimizations, sum optimizations, user-defined optimizations, and automated tradeoff.</p>
Comparing calibrations to data	<p>You can compare your calibrations to experimental data for validation.</p> <p>For example, after completing a calibration, you can import experimental data from a spreadsheet. You can use CAGE to compare your calibration to the data.</p>

Starting the CAGE Browser

To start the CAGE, at the MATLAB® command prompt, type:

cage

See Also

More About

- “Navigate CAGE” on page 1-4
- “Feature Calibration”
- “Tradeoff Calibration”

Navigate CAGE

CAGE includes a **Processes** pane and a **Data Objects** pane to help you identify the type of calibration you want to do and the data objects that you intend to use. Use the buttons in these panes to navigate between the different sections of functionality in CAGE. The view of CAGE depends on:

- Which button you select in the **Processes** and **Data Objects** panes.
- The item you highlight in the tree display.

Processes

The **Processes** pane has three buttons.

Button	Description
Feature	Shows the Feature view, with the tables and strategies that are associated with that feature. A feature is a strategy (or collection of tables) and a model used to calibrate those tables. In the Feature view, you can fill lookup tables by comparing a strategy to a model. You can import existing strategies or construct new ones using Simulink® software from the feature view. From the feature node in the tree display, you can access the Surface Viewer to examine the strategy or model or both.
Tradeoff	Shows the Tradeoff view, with a list of the tables and models to display. Here you can see graphically the effects of manually altering variables to trade off different objectives (such as maximizing torque while minimizing emissions). At the tradeoff node, you can calibrate table values to achieve the best compromise between competing objectives. You can calibrate using single or multimodel tradeoffs. You can also use the optimization functionality of CAGE to run automated tradeoffs.
Optimization	Shows the Optimization view. From here you can set up and run optimizations, including automated tradeoffs. There are standard routines available and also templates provided so you can write your own optimization routines.

Data Objects

The **Data Objects** pane has four buttons.

Button	Description
Variable Dictionary	Stores all the variables, constants, and formulas in your session. Here you can view, add, and edit any variables in any part of your session.

Button	Description
Models	Stores all the models in your session. Here you can view a graphical display of these models, including a diagram of the model input structure. This is useful because a model can have other models as inputs. You can change the inputs here. For example, you can change your model input Spark to be connected to a model for Spark rather than to the variable Spark. You can also access the surface viewer here to examine models.
Lookup Tables	Enables you to see all the lookup tables and normalizers in your session. You can also calibrate lookup tables manually here if you want. You can add and delete lookup tables from the project. From any table display (here, or in other views) you can access the History Display to manage changes in your lookup tables and normalizers. You can use the History Display to reverse changes.
Data Sets	Enables you to evaluate your models and features over a custom set of input values. Here you can create and edit a set of input values and view several models or features evaluated at these points. You can compare your tables and models with experimental data to validate your calibrations. You can also fill lookup tables directly from experimental data by loading the experimental data as a new data set.

See Also

More About

- “Feature Calibration”
- “Tradeoff Calibration”
- “What Is CAGE?” on page 1-2


Variables and Models

The following sections describe how to set up variables and models before performing calibrations.

Set Up Calibrations, Resume Work, or Find Calibration Examples

When you open MBC Optimization app, the home page helps you get started or resume work faster by accessing frequent tasks, recent projects, and featured examples.

- Get started by using the buttons in the right pane for common calibration tasks: import models; set up optimizations, lookup tables, feature strategies, or data sets; or export lookup tables.
- Resume work by opening projects from the **Recent projects** list.
- Open case study examples from the **Case Studies** list.
- View a summary of items in your project. To open other views, click Browse Project.
- From other calibration views, to return to the home page, click the Home toolbar button or select **File > Home**.



See Also

Related Examples

- “Import Models and Calibration Items Using CAGE Import Tool” on page 2-4
- “Create an Optimization” on page 7-6
- “Create Lookup Tables from a Model” on page 3-2
- “Set Up Designs and Models, Resume Work, or Find Engine Modeling Examples”

More About

- “Set Up Models” on page 2-11
- “Engine Calibration Applications”

Import Models and Calibration Items Using CAGE Import Tool

You can use the CAGE Import Tool to select items to import from any Model-Based Calibration Toolbox project file produced in CAGE or the Model Browser (.mat or .cag). This can greatly simplify setting up new projects, and also making changes to existing projects, for example to make use of new models in an existing optimization and calibration.


You can import Model Browser models from any project file or direct from the Model Browser when it is open. You can import the following CAGE items from any CAGE project: models (including feature and function models), variables, normalizers, lookup tables, features, optimizations, datasets and tradeoffs.

You can replace suitable items in your current CAGE project with imported items. You can see if an item is replaceable in the Import dialog, where the **Repl**ace action becomes available.

Note that Model Browser models (but not CAGE models) must have exactly the same input names as the CAGE model you want to replace. You can replace models, variables, normalizers, lookup tables and features. You cannot replace optimizations, datasets or tradeoffs. You cannot replace lookup tables used in tradeoffs with lookup tables of a different size.

To use the CAGE Import Tool:

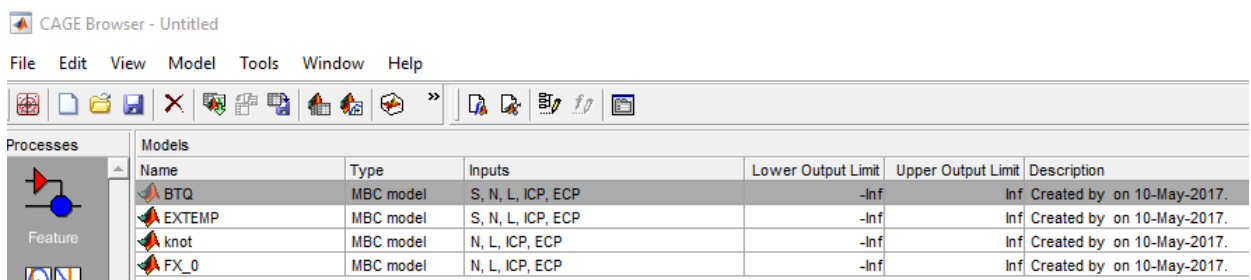
- 1** On the CAGE home page, click **Import Models**. Alternatively, select **File > Import From Project**.
 - If it is not already open, the Model Browser opens. Use the Model browser to open a project.
The CAGE Import Tool appears.
- 2** You can choose a project file or import directly from the Model Browser if it is open.
 - If the Model Browser is open, the list is automatically populated with a list of available items in the open project.
 - To import from a file, click the **Import from project file** button.
A file browser dialog opens. Locate the desired file and click **Open**.
- 3** The CAGE Import Tool displays the available items. Select the items you want to import from the list. Press **Ctrl+A** to select all items, or **Ctrl+click** or **Shift+click** to select multiple items in the list.



You can use the **Find** and **Type** controls to filter the item list:

- If you are importing from a Model Browser project you can select Response, Point-by-point, Datum or Response Feature from the **Type** list to display a single model type only.
- If you are importing from a CAGE project you can select Variable, Model, Normalizer, Table, Feature, Optimization, Dataset, or Tradeoff from the CAGE items in the **Type** list. For models the **Subtype** column displays whether a model item is an MBC model, function model or feature model.
- Enter text in the **Find** edit box to find particular item names. You can also select the box to **Match case**

- 4 Click the **Import Selected Items** button.
- 5 The CAGE Browser opens displaying the imported models.



- 6 When you have finished importing items, in the CAGE Import Tool, click **Close**.

See Also

More About

- “Importing and Exporting a Variable Dictionary” on page 2-9

- “Set Up Calibrations, Resume Work, or Find Calibration Examples” on page 2-2

Set Up Variables and Constants

In this section...

“Introducing the Variable Dictionary View” on page 2-7

“Importing and Exporting a Variable Dictionary” on page 2-9


“Using the Variable Menu” on page 2-9

“Using Aliases” on page 2-10

Introducing the Variable Dictionary View

The Variable Dictionary is a store for all the variables, constants, and formulae in your session.

To view or edit the items in the Variable Dictionary, click the button, shown, in the **Data Objects** pane.



Selecting the **Variable Dictionary** view displays the variables, constants, and formulae in the current project.

Following is an example of the **Variable Dictionary** view.

List of all the constants, variables, and formulas in the project

The screenshot shows the CAGE Browser software interface. The main window is titled "CAGE Browser - Untitled" and has a menu bar with "File", "Edit", "Variable", "Tools", "Window", and "Help". Below the menu bar is a toolbar with various icons. The interface is divided into several panes:

- Processes:** A vertical sidebar on the left containing icons for "Feature", "Manual Calibration", and "Tradeoff".
- Data Objects:** A vertical sidebar on the left containing icons for "Variable Dictionary", "Models", and "Data Sets".
- Variable Dictionary:** A table listing variables, constants, and formulas. The table has columns for Name, Type, Alias, Minimum, Maximum, Set Point, and Formula.
- Settings Pane:** A pane below the Variable Dictionary showing the settings for the selected variable "Afr". It includes fields for Alias, Description, Minimum, Maximum, Set Point, and Formula.

The Variable Dictionary table is as follows:

Name	Type	Alias	Minimum	Maximum	Set Point	Formula
x N	Variable	engine_speed	500	6500	2500	
x L	Variable	load, Load	0.1	1	0.4	
x A	Variable	afr, AFR	11	17	14.35	
k stoich	Constant				14.35	
x SPK	Variable	S, s, spark	-10	60	22.5	
f(x) lambda	Formula		0.75	1.25	1	A/stoich

The Settings Pane for the selected variable "Afr" shows the following values:

- Alias: afr, AFR
- Description: Air-fuel ratio (ratio)
- Minimum: 11
- Maximum: 17
- Set Point: 14.35
- Formula: (empty field)

Edit boxes to change the settings of the selected constant, variable, or formula

The upper pane shows a list of all the current variables, constants, and formulas. The lower pane displays edit boxes so you can specify the settings of the selected variable, constant, or formula.

Different Variable Dictionary Items

- Variables — standard items that feed into models, strategies and tables, and define ranges for these items
- Constant — used for inputs that you do not want to change
- Formulae — used when you want a variable item to depend on another

Importing and Exporting a Variable Dictionary

A variable dictionary contains all the variable items for your calibrations. You can set up your variable dictionary once, and use it in many calibrations.

If you import a model, it has variables associated with it, in which case you might not have to import a variable dictionary.

Importing a Variable Dictionary

To import a dictionary of variables from an .xml file,

- 1 Select **File > Import > Variable Dictionary**.
- 2 Select the correct dictionary file.

Exporting a Variable Dictionary

After setting up a variable dictionary, you can save the dictionary for use in many different calibrations.

To export a dictionary of variables to an .xml file,

- 1 Select **File > Export > Variable Dictionary**.
- 2 Select a suitable name for the dictionary file.

Using the Variable Menu

The **Variable** menu in the variable dictionary enables you to alter variable items. These choices are also available in the right-click context menu on the list view.

Change item to:

- **Alias**

Changes the selected item to be an alias of another item in the current project. For example, if you have two variables, `engine_speed` and `n`, you can change `n` to be an alias of `engine_speed`, with its maximum and minimum values.

- **Formula**

Changes a variable or constant into a formula. You have to define the right side of the formula, and you can select the check box to calculate the range.

- **Constant**

Changes a variable or formula into a constant. The value of the constant is the set point of the old item.

- **Variable**

Changes a constant or formula into a variable. The range is from 0 to twice the constant's value (negative values have a maximum of 0).

Using Aliases

The variable dictionary enables you to use the same set of variables, constants, and formulas with many different models and calibrations.

Why Use Aliases?

It is possible that in one model or strategy the engine speed has been defined as **N**, and in another it has been defined as **rpm**. The alias function enables you to automatically link inputs with various names to a single CAGE variable when you import models and strategies.

Creating an Alias

For example, in a variable dictionary there are two variables:

- **N**, with a range of 500 to 6500
- **rpm**, with a range of 2500 to 3500

To set **rpm** to be an alias of **N**,

- 1 Highlight the variable **rpm**.
- 2 Select **Variable > Change item to > Alias**.
- 3 In the dialog, choose **N** from the list.

This eliminates the variable **rpm** from your variable dictionary, and every model and calibration that refers to **rpm** now refers to **N** instead.

Note If **N** is made an alias of **rpm** in the preceding example, the range of **N** is restricted to the range of **rpm**, 2500 to 3500.

You can also add aliases to existing items by entering a list of names in the **Alias** edit box.

See Also

More About

- “Set Up Variables and Constants” on page 2-7
- “Set Up Calibrations, Resume Work, or Find Calibration Examples” on page 2-2
- “Using Aliases” on page 2-10

Set Up Models

In this section...

“Introducing the Models View” on page 2-11
 “Importing Models” on page 2-13
 “Adding New Function Models” on page 2-13
 “Renaming and Editing Models” on page 2-14

Introducing the Models View

CAGE generally calibrates lookup tables by reference to models.

To import models, see “Import Models and Calibration Items Using CAGE Import Tool” on page 2-4.

The **Models** view is a storage place for all the models in your session.

To view and edit the models in your session, select **Models** by clicking the button shown in the **Data Objects** pane.




The **Models** view displays the following:

- A list of all the models in the current project.
- The model connections. That is, which constants, variables, and models are inputs to the selected model. You can use the **View** menu or the right-click context menu on the graph to zoom in and out, zoom to fit, and reset.
- An image of the response surface of the selected model; you can select factors to display. Use the **View** menu to choose between:
 - **No Constraint Display** — Shows entire model surface.
 - **Show Constraint** — Areas outside the boundary constraint model (if any) are yellow.
 - **Clip to Constraint** — The surface is only shown within the boundary constraint model.

View > Edit Input Set Points opens a dialog box where you can edit the set points of your model variables. This setting alters the model display and also any calculations involving the set points throughout CAGE.

Following is an example of the **Models** display.



The icons in the Models list indicate the type of model, as listed in the Type column. As shown in the following illustration, a model can be a Model Browser statistical model, the boundary of a model, the prediction error variance (PEV) of a model, a user-defined function model, or a feature model (converted from a feature).

-  Stats_model
-  Boundary_model
-  Function_model
-  PEV_model
-  Feature_model

Use the Model Properties dialog box to switch a model output between the model value and the boundary or PEV of the model. For function models see “Adding New Function Models” on page 2-13. You can convert a feature to a model by selecting **Feature > Convert to Model**.

Importing Models

CAGE enables you to calibrate lookup tables by referring to models constructed in the Model Browser.

CAGE can only open Model-Based Calibration Toolbox model files. You can import models from project files (.mat, .cag).

Import Models From Project

You can use the CAGE Import Tool to select models to import from any Model-Based Calibration Toolbox project file produced in CAGE or the Model Browser (.mat or .cag). You can replace suitable models in your current CAGE project (note that Model Browser models must have exactly the same input names as the CAGE model you are replacing).

Adding New Function Models

A function model is a model that is expressed algebraically. The function can be any MATLAB function (including user-defined functions). The only restriction is that the function must be vectorized, that is, take in column vectors and return a column vector of the same size, as in this example:

```
function y = foo(x1, x2)
y = x1 .* x2;
```

Once you have a function like this, you can create a function model applying it to any models or variables in your session, like the following example.

```
foo(NOX, SPK)
```


For example, you might want to view the behavior of torque efficiency. So you create a function model of torque efficiency = torque/peak torque.

To add a function model to your session,

- 1 Select **File > New > Function Model**.

This opens the Function Model Wizard.

- 2 In the dialog box, enter the formula for your function model. For example, enter `torque_efficiency=torque/peak_torque`.
- 3 Press **Enter**. CAGE checks that the function is recognized; if so, you can click **Next**. If the function is incorrectly entered, you cannot click **Next**.
- 4 Select the models that you want to import by highlighting the models from the list.
- 5 Click **Next**.



- 6 You can select the check box to **Automatically assign/create inputs** and click **Finish** to close the wizard and return you to the **Models** view, or you can click **Next** and go to the next screen. Here you can manually associate the model factors with the available inputs as follows:
 - a Highlight a **Model Input**, e.g., `peak_torque`, in the list on the left and the corresponding model, `peak_torque`, in the **Available Inputs** list on the right.
 - b Click the **Assign input** button.

Repeat a and b for all the model factors. Click **Finish** to close the wizard and return you to the **Models** view.

You can now see a display of the model and its connections (inputs).

Renaming and Editing Models

Renaming Models

To rename a model,

- 1 Highlight the model that you want to rename.
- 2 Select **Edit > Rename**.
- 3 Enter the new name for the model and press **Enter**.

You can also rename the model by selecting a model and clicking the name, or pressing **F2**.

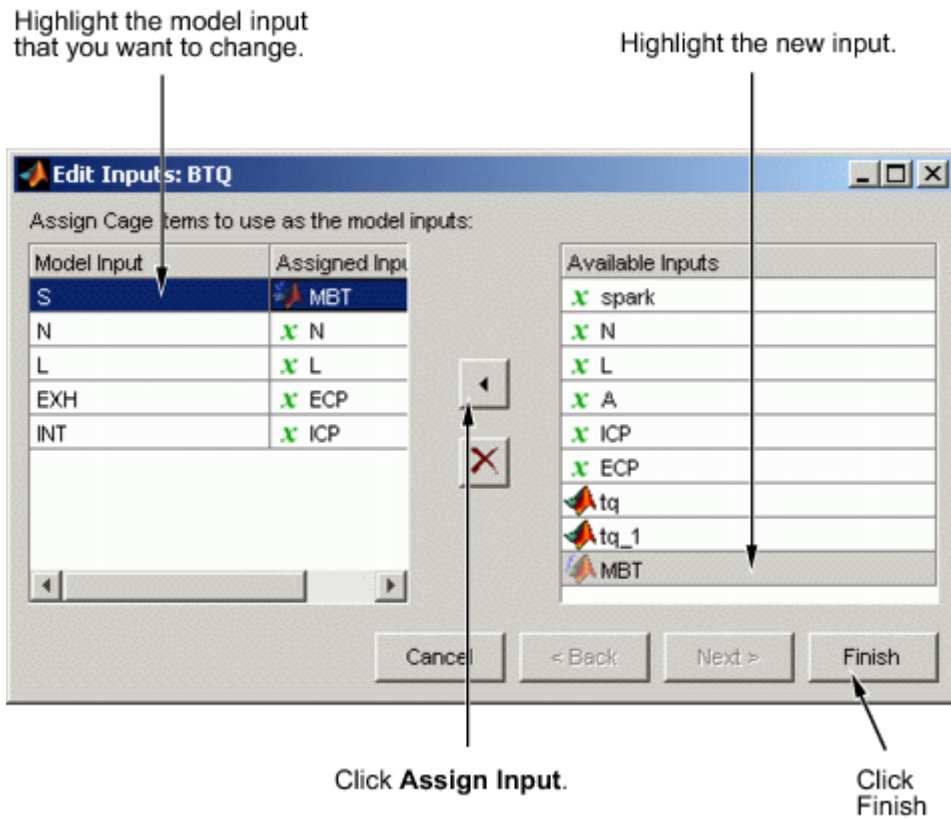
Editing Model Inputs

You can adjust a model so that variables, formulas, or other models are the factors of the model. For example, a model of torque depends on the spark angle. In place of the spark angle variable, you can use a model of the maximum brake torque (MBT) as the spark input.

To edit the inputs of a model,

- 1 Highlight the model.
- 2 Select **Model > Edit Inputs**.

This opens the Edit Inputs dialog box, shown.



- 3 Highlight the **Model Input** that you want to edit, in the list on the left.
- 4 Highlight the new input for that factor, in the **Available Inputs** list on the right.
- 5 Click the **Assign Input** button.
- 6 To close the dialog box, click **Finish**.

Note If you want to change the range of a variable in the session, change the range in the variable dictionary.

See Also

More About

- “Import Models and Calibration Items Using CAGE Import Tool” on page 2-4
- “Set Up Calibrations, Resume Work, or Find Calibration Examples” on page 2-2
- “Using the Variable Menu” on page 2-9

Composite Models

In this section...
“What Are Composite Models?” on page 2-16
“Importing from the Model Browser” on page 2-16
“Combining Existing CAGE Models” on page 2-19
“Viewing Composite Model Properties” on page 2-20

What Are Composite Models?

The composite model type allows you to combine a number of models to represent engine responses under different operating modes. You can use the composite model in CAGE to produce optimal calibrations for engines with multiple operating modes. Use composite models for calibration problems where the goal is to fill a single table for all modes or to fill a table for each mode, such as:

- Multi-injection diesel engine
- Inclusion of startup conditions in drive cycles
- Rich and stoich regions for engines

You can create a composite model using either of the following approaches:

- Importing into CAGE from the Model Browser and combining suitable responses from different test plans
- Combining existing models in CAGE

The composite model comprises a collection of models with an extra mode input. The mode input is an index into the list of models. The component models can have different sets of inputs. CAGE uses symbol names and not signal names to determine inputs. Make sure the symbol names match for the same inputs.


Importing from the Model Browser

To import and combine models in the Model Browser into a single composite model:

- 1** Open the Model Browser, and load the project you want to import composite models from.
- 2** In CAGE, select **File > Import > Composite Model**.


The Import Composite Model wizard appears.

- 3** Select the test plans from which you want to import models, and then click **Next**.



- 4 Select the responses to combine into a composite model. You can only select responses that are common to all selected test plans. If you do not want a boundary model included, clear the **Include Boundary** check box.


Note Using composite models can result in very large CAGE project files when you combine point-by-point test plans with a large number of responses. To reduce the size of project files, exclude boundary models from response models that you do not want to use as optimization objectives.



After making your selections, click **Next**.

- 5 Optional — Edit the name for your mode variable and the names of your operating modes.

If you have point-by-point models, you can edit the input tolerances to merge close operating points.



- 6 Click **Finish** to import your new composite model.


Combining Existing CAGE Models

To combine existing CAGE models into a single composite model:

- 1 Select **File > New > Composite Model** (or use the toolbar button).

The New Composite Model wizard appears.

- 2 Select the models to combine into a composite model. The combined inputs are listed for your selected models. You can combine additional models with existing composite models as needed. Click **Next**.



- Optional — Edit the name for your mode variable and the names of your operating modes.


If you have point-by-point models, you can edit the input tolerances to merge close operating points.

- Click **Finish** to create your new composite model.


Viewing Composite Model Properties

Select a composite model in the Models view, and select **Model > Properties**.

In the Composite Model Properties dialog box, click the **Model** tab to view information about the model modes and inputs, as the following figure shows.



For composite models created by combining point-by-point models, use the Composite Model Properties dialog box to view which modes are available for each operating point, as shown in the next figure.



See Also


More About

- “Set Up Models” on page 2-11
- “Importing Models” on page 2-13
- “Renaming and Editing Models” on page 2-14

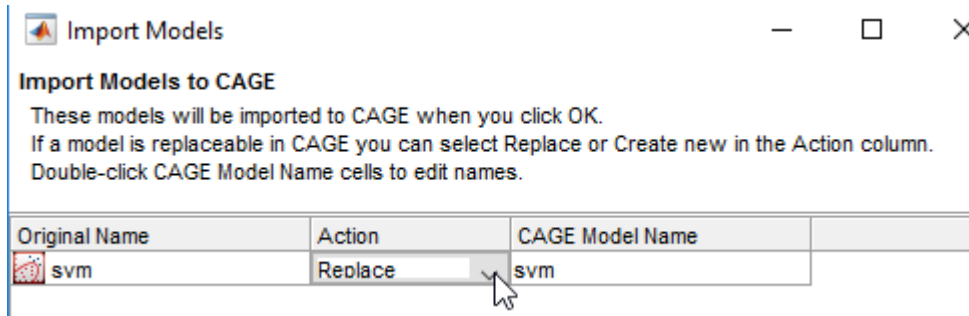
Classification Models

Using classification models, including support vector machine (SVM) or discriminant classifiers, Model-Based Calibration Toolbox can differentiate between operating modes when you specify boundary constraints for calibrations.

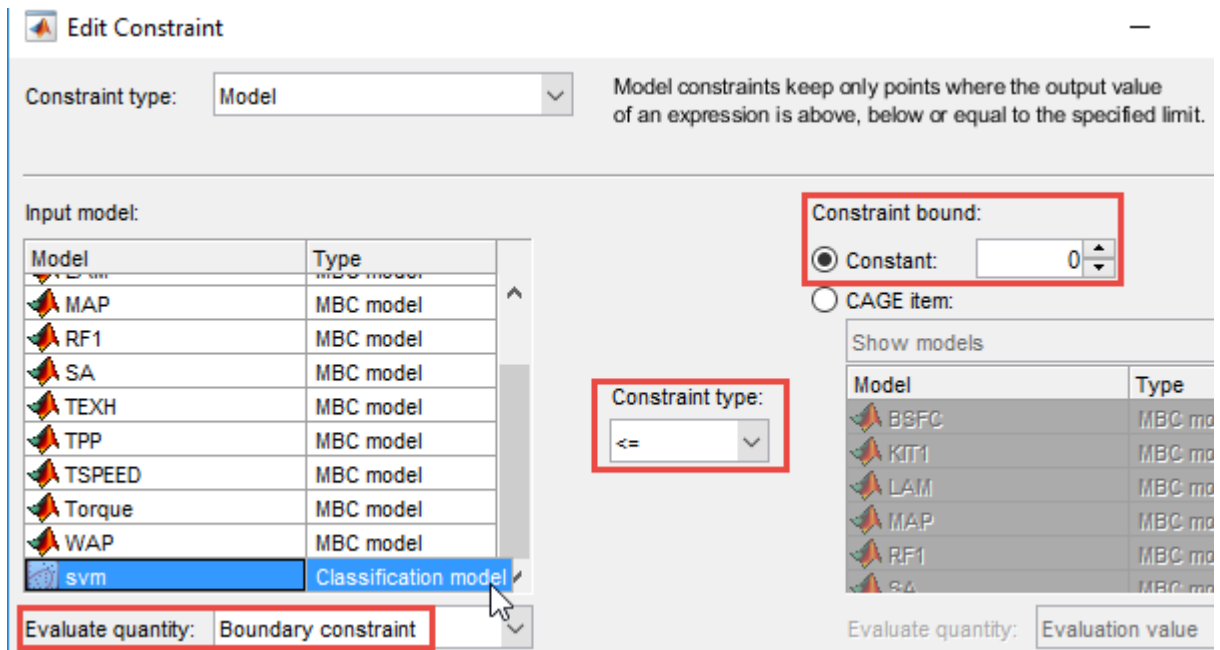
- 1 Use the Statistics and Machine Learning Toolbox™ to classify data using a classification model, for example a discriminant or SVM model.
- 2 Use the “Classification Learner App” to export the classification model to the workspace.
- 3 In CAGE, select **File > Import > Classification Model**.
- 4 In the Import dialog box, click **OK** to import the classification model, for example svm.



- 5 In the Import Models dialog box, accept the defaults and click **OK**.



- 6 In CAGE, select **Optimization > Constraints > Edit Constraint**. For the svm input model:
 - Set **Evaluate quantity** to Boundary constraint.
 - Specify the **Constraint type** and **Constraint bound**.



See Also

More About

- “Classification Learner App”
- “Support Vector Machine Classification”

Regression Models

Regression models describe the relationship between a response (output) variable, and one or more predictor (input) variables. Statistics and Machine Learning Toolbox allows you to fit linear, generalized linear, and nonlinear regression models, including stepwise models and mixed-effects models. Once you fit a model, you can use it to predict or simulate responses, assess the model fit using hypothesis tests, or use plots to visualize diagnostics, residuals, and interaction effects.

You can use CAGE to import and optimize Statistics and Machine Learning Toolbox regression models.

- 1 Use the Statistics and Machine Learning Toolbox to create a regression model.
- 2 Use the “Regression Learner App” to export the regression model to the workspace.
- 3 In CAGE, select **File > Import > Regression Model**.
- 4 In the Import dialog box, click **OK** to import the regression model.
- 5 In the Import Models dialog box, accept the defaults and click **OK**.
- 6 Use CAGE to optimize the regression model.

See Also

More About

- “Set Up Models” on page 2-11
- “Regression”
- “Regression Learner App”

Tables

This section includes the following topics:

- “Create Lookup Tables from a Model” on page 3-2
- “Import Lookup Tables from a Simulink Model” on page 3-5
- “Edit Lookup Tables” on page 3-7
- “Fill a Single Lookup Table From a Model” on page 3-14
- “Compare and Revert Lookup Tables” on page 3-17
- “Calibration Manager” on page 3-20
- “Lookup Table Properties” on page 3-24
- “Lookup Table Normalizers” on page 3-29
- “Invert Lookup Tables” on page 3-35
- “Import and Export Calibrations” on page 3-39

Create Lookup Tables from a Model

You can access the lookup table creation wizard by menu or toolbar, from any view in CAGE. The wizard helps you quickly create a set of lookup tables with the same axes for all the inputs of a model, and the model response, and any other responses that share inputs. You can choose which of these lookup tables to create, and select the values for the axes (or normalizers) that all lookup tables share. You can also add all the new tables to a tradeoff.

This wizard can be useful when creating tables for an optimization, to use when filling tables with optimization results, and for investigating results in the tradeoff views.

To create lookup tables (and optionally a tradeoff) from a model,

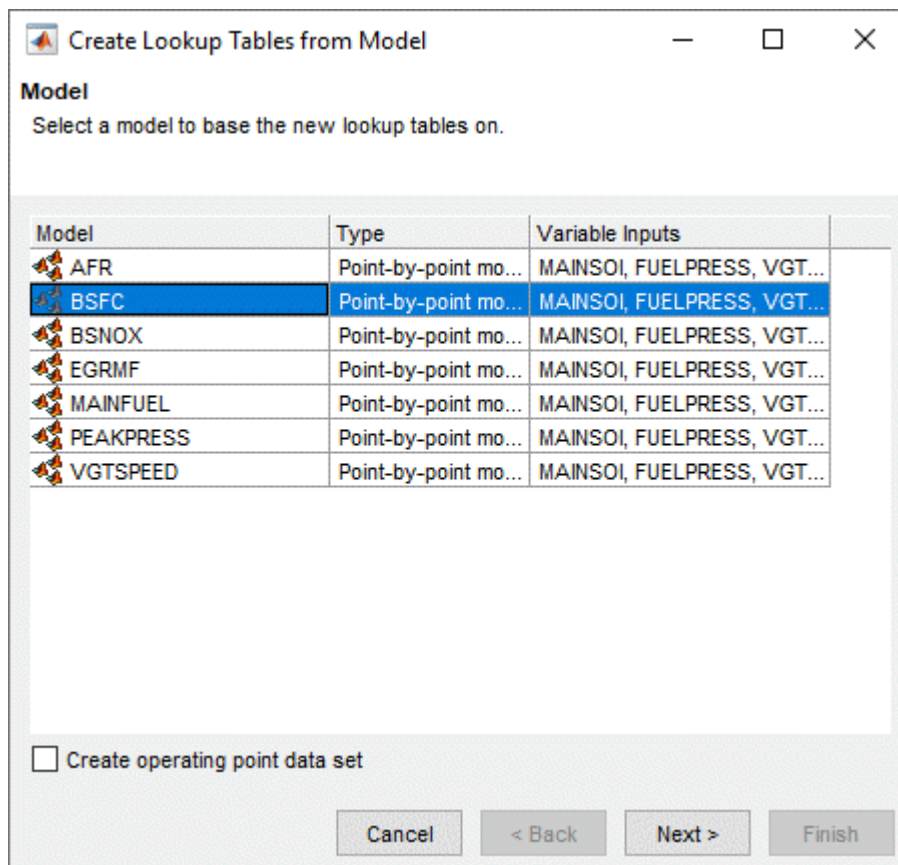
- 1 Select **Tools > Create Lookup Tables From Model**.

The **Create Lookup Tables From Model** Wizard opens.

- 2 Select a model to base the new tables on.

If you are viewing a model, then the wizard automatically selects the current model. If you are viewing an optimization or an optimization output node, then the wizard automatically selects the model in the first objective. You can use this to create lookup tables for the selected optimization.

If you have selected a point-by-point model, you can optionally select the check box to **Create operating point data set**.



Click **Next**.

- 3 Select table axes input variables and set up the normalizers to use for the new tables.

Lookup Table Inputs
Select the lookup table inputs and set up the normalizers to use for all the new lookup tables.

Use model operating points

Rows (Y) input: SPEED
Normalizer: <New>
Table rows: 7

Columns (X) input: BTQ
Normalizer: <New>
Table columns: 7

SPEED normalizer:

Input	Output
1600	0
1700	1
1800	2
1900	3
2000	4
2100	5

BTQ normalizer:

Input	Output
0	0
266.667	1
533.333	2
800	3
1066.667	4
1333.333	5

Buttons: Cancel, < Back, Next >, Finish

If you have selected a point-by-point model, CAGE automatically selects the check box to **Use model operating points** for the table normalizers. You can clear the check box if you want to select different normalizers.

- Select inputs. For the **X-** and **Y-axis inputs**, you can select any input variable for your selected model, or the model response.
- Select normalizers. You can select existing normalizers in your project or create ones. If creating normalizers you can edit the numbers of **Table rows** and **Table columns**, and edit values in the **Input** columns. By default CAGE initializes normalizers with equally spaced points across variable ranges, unless you select the response model as a table input.

If you choose a response model input, you must specify the breakpoints. Click the button to **Edit breakpoints**, then enter a number of points and the range to space the breakpoints over. If you do not do this, model inputs are spaced over 0-1, because CAGE cannot determine the range automatically as happens with variables. After you create your tables with a model input, in your Variable Dictionary you can view a new variable named `modelname_input` with the range you specified. CAGE uses this input variable to match to model names when you fill lookup tables from optimization results.

Click **Next**.

- 4 Select check boxes to specify which variables and responses to create lookup tables for. You can create lookup tables for other responses with the same inputs as the primary model (and the same operating points for point-by-point models).

By default, you create a tradeoff containing all of the new tables. The tradeoff can be useful for investigating optimization results. If you do not want to create the tradeoff, clear the check box.

Click **Finish** to create the tables and tradeoff.

You see a dialog box listing all the items you have created, and you can choose which (if any) of the items to view next.

See Also

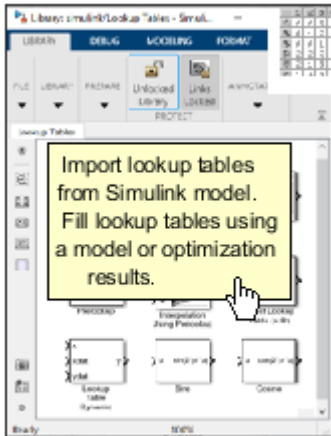
More About

- “Edit Lookup Tables” on page 3-7
- “Lookup Table Filling When Optimization Operating Point Inputs Differ from Table Inputs” on page 8-11

Import Lookup Tables from a Simulink Model

On the CAGE browser home page, use Simulink Lookup Tables to import 1-D and 2-D lookup tables from a Simulink model. After you import the lookup tables, you can fill them using statistical models or optimization results.

- 1 In the CAGE browser home page, select **Simulink Lookup Tables**.



Simulink Lookup Tables

Use the browser to open the Simulink model with the lookup tables that you want to import.

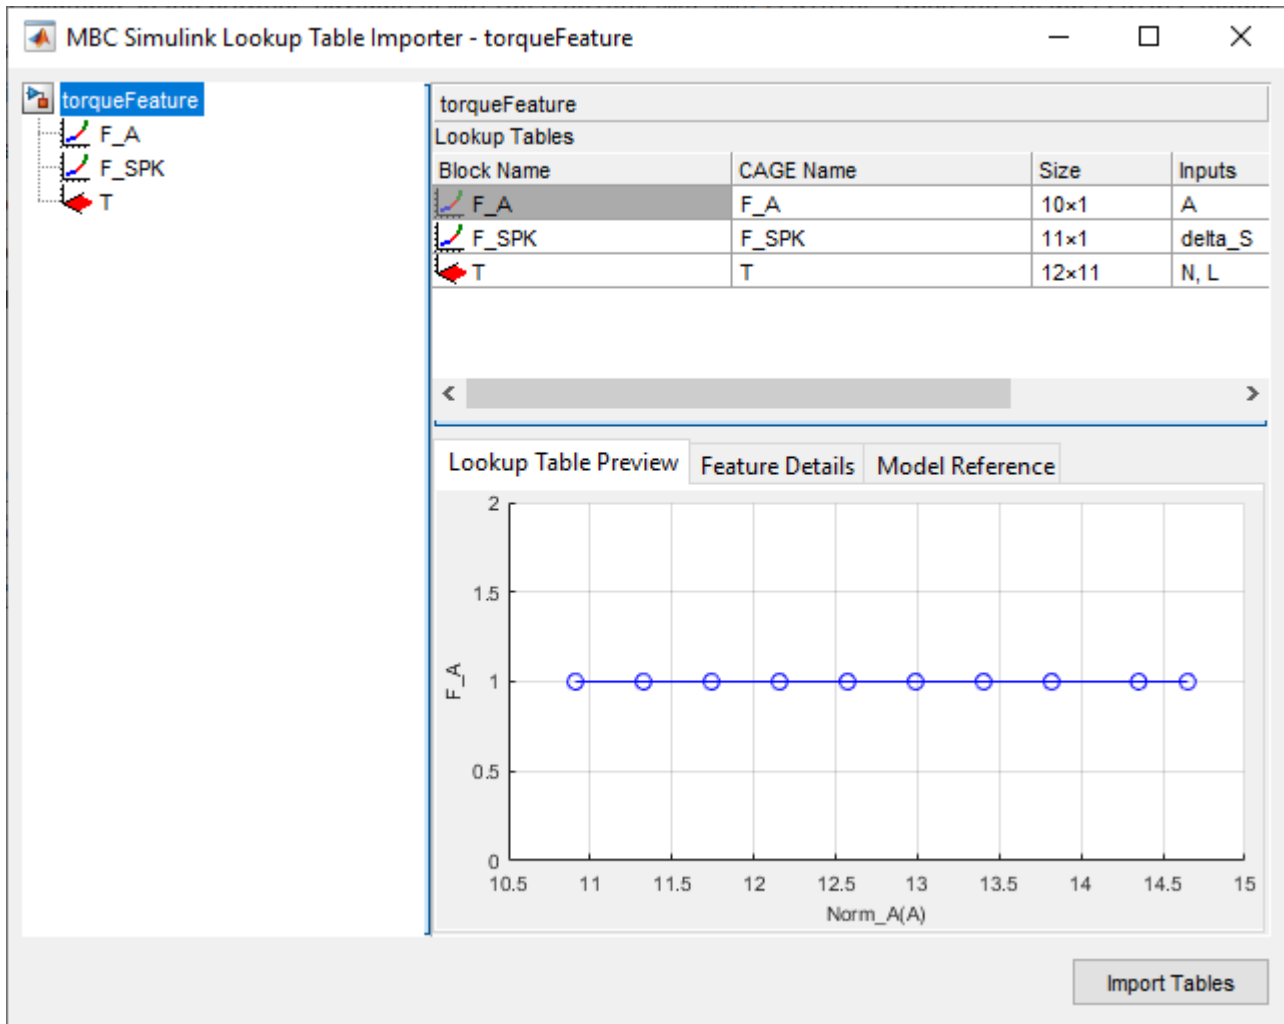
For example, in the browser, navigate to `matlab\toolbox\mbc\mbctraining`. Open the `torqueFeature` model.

- 2 Use the MBC Simulink Lookup Table Importer to preview and import the 1-D and 2-D lookup tables in the model, including lookup tables in subsystems.

Click **Import Tables** to import tables into CAGE.

Tabs	Description
Lookup Table Preview	Preview of the table that you select in the Lookup Tables pane.
Feature Details	The inputs and outputs for subsystems in the model.
Model Reference	List of model references in the model. To import lookup tables in a model reference, double-click the model. The model reference opens. The table importer contains the model reference lookup tables.

For example, the MBC Simulink Lookup Table Importer provides this information for the `torqueFeature` model lookup tables.



Note The MBC Simulink Lookup Table Importer remains opened until you close CAGE or open another project.

- In CAGE, you can fill the lookup tables using statistical models or optimization results.

See Also

More About

- “Edit Lookup Tables” on page 3-7
- “Fill a Single Lookup Table From a Model” on page 3-14

Edit Lookup Tables

In this section...
“Add, Duplicate, and Delete Tables” on page 3-8
“Fill Tables by Extrapolation” on page 3-9
“Optimize Lookup Table Breakpoints” on page 3-11
“Lock Lookup Table Values” on page 3-12
“Edit the Lookup Table Graph” on page 3-12
“Arithmetic Operations on Lookup Table Values” on page 3-12

When you select a table in the tree (under feature or tables), you see the **Lookup Tables** view.

In CAGE, a lookup table is defined to be either a one-dimensional or a two-dimensional lookup table. One-dimensional tables are sometimes known as characteristic lines or functions. Two-dimensional tables are also known as characteristic maps or tables. CAGE regards them both as similar objects.

Each lookup table has either one or two axes associated with it. These axes are normalizers.

For example, a simple MBT feature has two lookup tables:

- A two-dimensional lookup table with speed and relative air charge as its normalizers
- A one-dimensional lookup table with AFR as its normalizer

L \ N	500	1000	1500	2000	2500	3000
0.1	50	50	50	50	50	50
0.2	43.66	47.336	50	50	50	50
0.3	34.061	37.626	41.838	43.348	44.735	46.39
0.4	23.821	25.32	27.091	28.708	30.098	32.209
0.5	15.301	16.263	17.685	19.995	21.746	23.958
0.6	9.132	10.781	13.202	16.284	18.334	20.24
0.7	4.486	6.934	10.052	13.478	16.314	18.242
0.8	0.509	3.518	7.028	10.776	14.307	17.329
0.9	0	0.027	3.742	7.633	11.387	14.601
1	0	0	9.342e-3	3.85	7.506	10.665

In the Lookup Table view:

To	Action
Fill the table.	On the Common Tasks pane, click Fill . For an example, see “Fill a Single Lookup Table From a Model” on page 3-14.
Optimize the lookup table breakpoints.	On the Common Tasks pane, click Breakpoints Fill .
Extrapolate the data.	On the Common Tasks pane, click Extrapolate .
Edit the table properties.	On the Common Tasks pane, click Properties .
Interactively change the smoothing factor when you fill lookup tables with models or data.	Use the Lookup Table Smoothing interactive control to adjust the smoothing factor.
Edit a value in the lookup table.	Double-click the cell, then enter a value. You can right-click to Copy or Paste values.
Help you relate the lookup table values to surface plots, the lookup table is shaded. Table cells background color matches the corresponding surface plot color.	Toggle lookup table shading with View > Shade Table .
Match the orientation of the numerical lookup table.	Transpose the plot axes by right-clicking and selecting 2D Orientation X/Y . Switch back with 2D Orientation Y/X .
View and revert table changes.	Click View History . Any locked cells and extrapolation mask cells are saved with lookup table history.

The table displays the values of your lookup table and displays the breakpoints of the normalizers.


The table breakpoint values are not necessarily identical to the normalizer breakpoints. When you create a table, the breakpoint values are the same as the normalizer values. If you delete breakpoints from the normalizers the table size does not change, so the table column and row breakpoint values are interpolated between the remaining normalizer breakpoints.

Add, Duplicate, and Delete Tables

To quickly create lookup tables from a model, use the lookup table creation wizard. To add tables, select the **Lookup Tables** view, or CAGE automatically switches to this view if you add a table using the **File > New** menu items.




The **Lookup Tables** view lists all the tables and normalizers in the current CAGE session.

To	Action
Add a lookup table to a session:	<p>1 Decide whether you want to add a one- or a two-dimensional table.</p> <p>For example if you want to add a modifier table to account for the variation in exhaust gas recirculation, add a one-dimensional table (which has one input). If, however, you want to add a table with speed and load as its normalizer inputs, then add a two-dimensional table.</p> <p>2 Select File > New > 1D Lookup Table or File > New > 2D Lookup Table as appropriate.</p> <p>Adding new tables automatically switches you to the Lookup Tables view.</p> <p>3 In the Create Lookup Table dialog box you can specify the table:</p> <ul style="list-style-type: none"> • Name • Range • Rows • Normalizer • Breakpoints with the Edit breakpoints option <p>4 Click OK to add the new table.</p>
Copy a table or a normalizer from a session.	<p>1 Select the Lookup Tables view.</p> <p>2 Highlight the required table or normalizer.</p> <p>3 Select Edit > Duplicate <i>table_name</i> (<i>table_name</i>' is the currently selected table).</p>
<p>Delete a table or a normalizer from a session.</p> <p>When you are calibrating a collection of lookup tables using either Feature or Tradeoff calibrations, you cannot easily delete tables without affecting the entire calibration. When deleting items, you must delete from the highest level down. For example, you cannot delete a table that is part of a feature; you must delete the feature first.</p>	<p>1 Select Lookup Tables view.</p> <p>2 Highlight the required lookup table or normalizer.</p> <p>3 Click ; or press Delete; or select Edit > Delete <i>table_name</i> (<i>table_name</i>' is the currently selected table).</p>

Fill Tables by Extrapolation

Filling a table by extrapolation fills the table with values based on the values already placed in the extrapolation mask.

To fill a table by extrapolating over a preselected mask, click  or select **Table > Extrapolate**. Alternatively, on the **Common Tasks** pane, click **Extrapolate**.

This extrapolation does one of the following.

If Extrapolation Mask Has	Action
Points on a line.	CAGE performs linear extrapolation on points projected on to that line. The simplest case of this is when you try to fill a 2-D table using data from a single row or column.
Points on a plane.	CAGE uses the plane for extrapolation. The simplest example is when the mask has three points and the points are not on a line.
Four or more ordered cells in a grid	CAGE uses bilinear extrapolation.
Four or more cells not on a grid,	CAGE uses a thin plate spline to extrapolate the table values.

To ensure that tables extrapolate smoothly to locked cell values, CAGE includes the locked cells in the extrapolation even if the cells are not part of the extrapolation mask.

Extrapolation Masks

The extrapolation mask defines a set of cells that form the basis of any extrapolation. For example, a speed-load (or relative air charge) table has values in the following ranges that you consider to be accurate:

- Speed 3000-5000 rpm
- Load 0.4-0.6

To add or remove a cell from the extrapolation mask:

- 1 Right-click the table.
- 2 Select **Add To Extrapolation Mask** or **Remove From Extrapolation Mask** from the menu.

Cells included in the extrapolation mask are colored blue. Cells that are locked and in the extrapolation mask are blue and have a padlock icon.

Use the **Table** menu for these options:

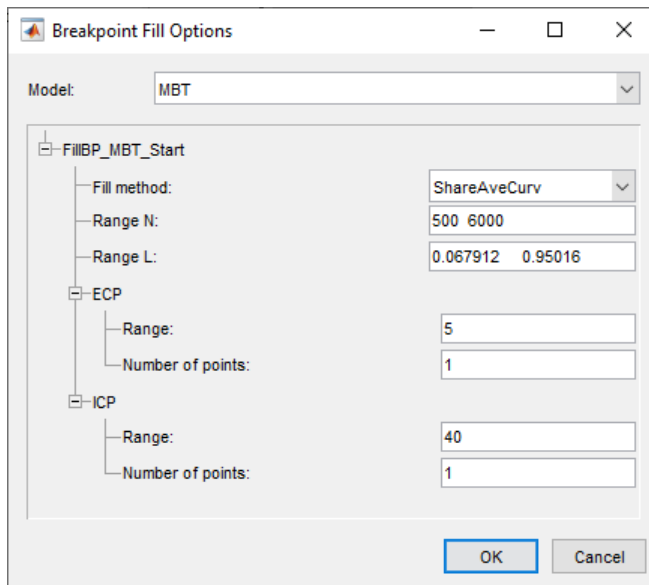
Menu	Action
Table > Extrapolation Mask	Use options to add, remove, clear, copy, and convert masks.
Table > Extrapolation Mask > Generate From PEV	Generates extrapolation mask depending on the value of prediction error (PEV). A dialog box opens where you can specify the threshold value of PEV below which you want to include cells in the mask. To help you select a threshold, the dialog box contains information about the range and prediction error mean for the model.
Table > Extrapolation Mask > Generate From Boundary Model	Generates extrapolation mask to include only cells within the boundary model.

Menu	Action
Table > Extrapolation	Extrapolates values from the cells in the extrapolation mask to fill the whole table.

Optimize Lookup Table Breakpoints

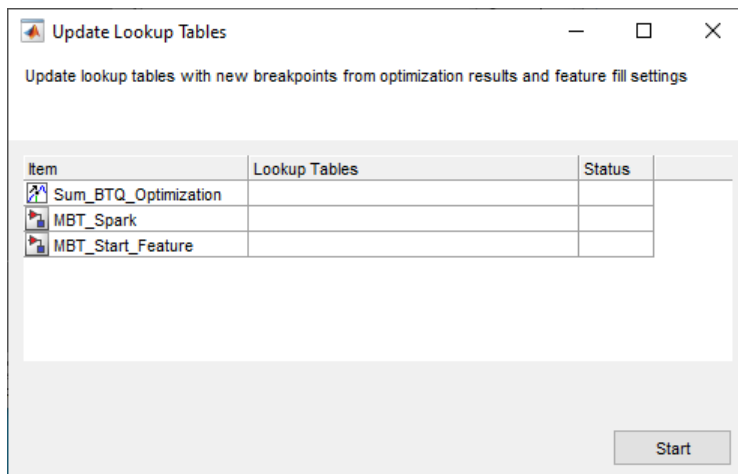
- 1 On the Lookup Table view **Common Tasks** pane, select **Breakpoints Fill**. On Breakpoint Fill Options dialog box, use the **Model** parameter to specify the model that you want to use for the optimization. Specify the other fill options.

Click **OK**.

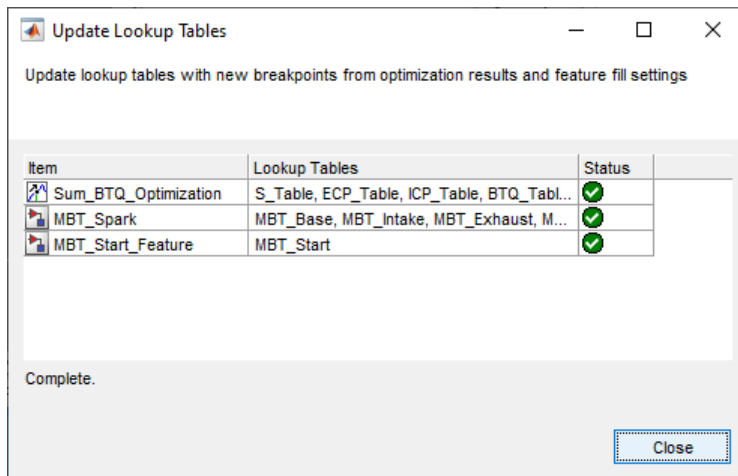


To help you update the lookup tables with the new breakpoints from the optimization results, use the Update Lookup Tables.

- 2 On the Update Lookup Tables dialog box, update the tables with new breakpoints from the optimization results and feature fill settings. Click **Start**.



The dialog box provides the status of the lookup table updates.



Lock Lookup Table Values

When you are satisfied with an area of the table, you might want to lock the cell values, to ensure that those values do not change.

To lock or unlock a cell value, right-click the cell and select from the menu. Locked cells have a padlock icon in the display. You can also lock an entire table or copy locks from other tables.

Use the **Table > Table Cell Locks** menu options to lock, unlock, copy, or convert selected or all the cells in the table.

Edit the Lookup Table Graph

The Lookup Tables view displays both the lookup table values and a graph of the lookup table. This gives a useful display of the table behavior. Use the **View** options to edit and alter the table views.

Arithmetic Operations on Lookup Table Values

The **Table > Adjust Cell Values** menu option opens a dialog box where you can specify an arithmetic operation to apply to either the whole table or only the cells currently selected. Arguments to operations can be numeric (plus 10) or percentages (minus 5%). You can set the selected cells to a value or to the mean. You can also apply user-defined functions.

- 1 Right-click the table or select **Table > Adjust Cell Values**. The Adjust Cell Values dialog box appears.
- 2 Select the operation to apply. Use the custom operation to specify your own function in a file.
- 3 Use the **Value** edit box to enter an argument. All operators accept a numeric argument (e.g. operator = plus, value = 10). You can also enter a percentage for the operators plus, minus, and set to value (e.g. 'minus' `1%').
- 4 Select the radio buttons to apply the operation to either the whole table or only the cells currently selected, and click **OK**.

You can use the custom operation option to apply user-defined functions. For example, use this to call the custom function:

```
newvalues = customfcn( currentvalue, selectedregion )
```

- `currentvalue` — Matrix of table values and `selectedregion` is a logical matrix the same size as the table.
- `newvalues` — Matrix should be the same size as `currentvalue`, and these numbers are put straight into the table.

For example:

```
function table = addOne( table, region )
table(region) = table(region) + 1;
return;
```

```
function table = randomtable( table, region )
table( region ) = rand( nnz( region ), 1 );
```

```
function table = saturate( table, region )
maxValueAllowed = 150;
table( region & table>maxValueAllowed ) = maxValueAllowed;
minValueAllowed = 100;
table( region & table<minValueAllowed ) = minValueAllowed ;
return
```

To use the `saturate` example:

- 1 Save the function text in a file named `saturate.m`.
- 2 Click and drag to select a region of cells in a CAGE table.
- 3 Right-click and select **Adjust Cell Values**.
- 4 In the dialog box:
 - Select `custom operation` from the **Operation** list
 - Enter `saturate` in the **Value** edit box or click the browse button to locate the file.
 - Select the option button to **Apply to selected table cells**, and click **OK**.

The selected table cells are saturated between the ranges specified in the function file (between 100-150).


See Also

More About

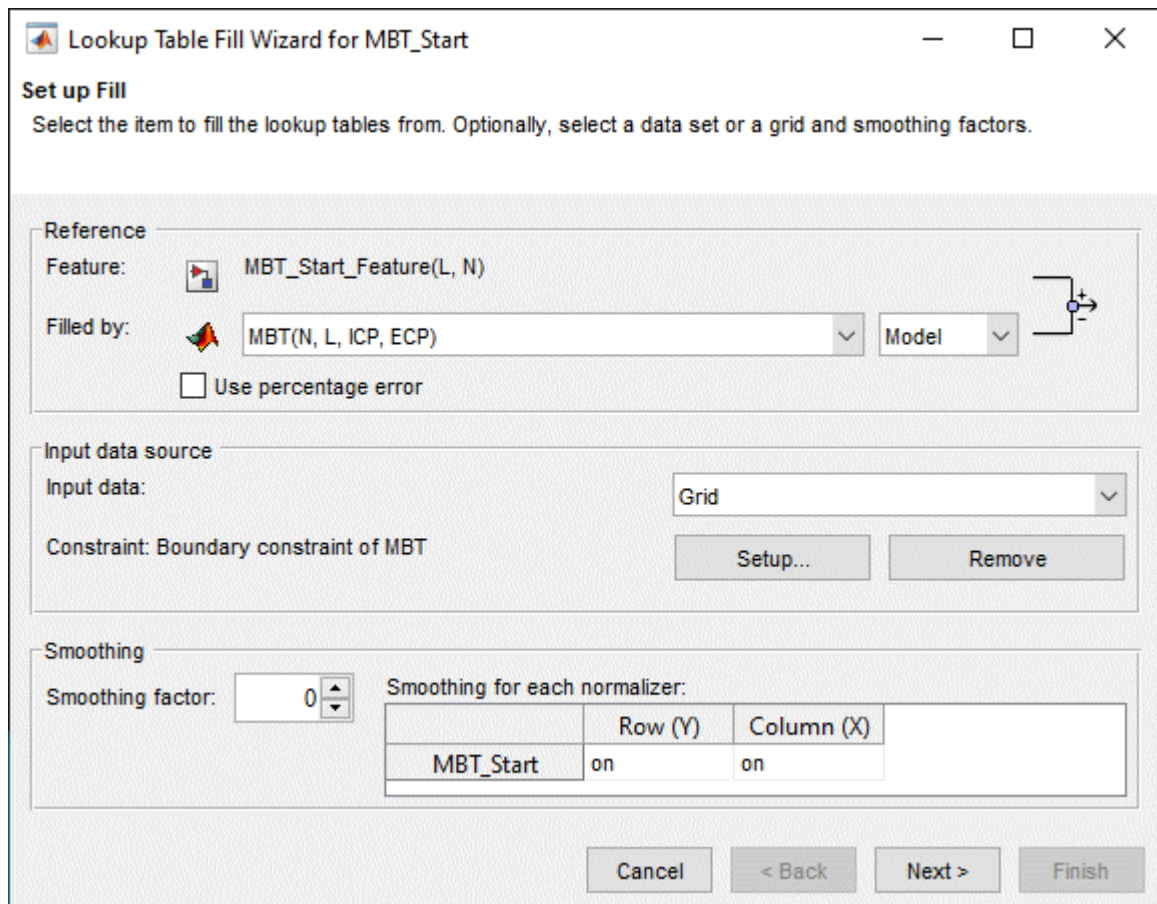
- “Fill Lookup Tables from Data”
- “Lookup Table Properties” on page 3-24
- “Lookup Table Normalizers” on page 3-29

Fill a Single Lookup Table From a Model


To fill a table with values from a model evaluated at the table breakpoints:

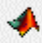
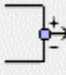
- 1 Open the Lookup Table Fill Wizard.
 - With the table selected in the Lookup Tables view, select **Table > Fill**.
 - Alternatively, click the Fill toolbar button ().

The Lookup Table Fill Wizard opens.



Reference

Feature:  MBT_Start_Feature(L, N)

Filled by:  MBT(N, L, ICP, ECP) Model 

Use percentage error

Input data source

Input data: Grid

Constraint: Boundary constraint of MBT Setup... Remove

Smoothing

Smoothing factor: 0

Smoothing for each normalizer:

	Row (Y)	Column (X)
MBT_Start	on	on

Cancel < Back Next > Finish

- 2 Use the **Filled by** options to select the model that you want to fill the table with.


- 3 (Optional) Specify the **Input data**, **Constraint**, and **Smoothing factor**. Click **Next**.

	Row (Y)	Column (X)
MBT_Start	on	on

- 4 Connect model and feature inputs. Eliminate inputs by connecting inputs to tables. Connect variables that describe the same quantities. Click **Next**.
- 5 Set the input grid values that you want to use for the feature fill. Click **Finish**.

The progress graph shows the change in RMSE as the optimization progresses.

- 6 On the Lookup Tables view, use the **Smoothing factor** to interactively change the smoothing factor when you fill lookup tables.



See Also

More About

- “Initialize Feature and Calibration Tables” on page 4-12
- “Optimize Lookup Table Breakpoints” on page 3-11

Compare and Revert Lookup Tables

In this section...

“Reset to Previous Lookup Table Versions” on page 3-17

“Compare Versions” on page 3-18

The History display enables you to view the history of any table or normalizer in a CAGE session. You can:

- Revert to previous versions of tables and normalizers. Any locked cells and extrapolation mask cells are saved with table history. Reverting to a version includes reverting to that version masks and locks (except any table history from release earlier than R2014a).
- Compare different versions of tables and normalizers.

You can view the History display of a table or normalizer by selecting **View > History**, or clicking **View History** in the Table Details pane.

Version	Comment / Action	Date and Time
3	Extrapolated	29-Nov-2013 16:44:46
2	Values filled from optimization output Sum_BTQ_Optimization_Output ...	30-Apr-2009 13:26:01
1	Table for 'S'	30-Apr-2009 11:55:38

N	L	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
500		50	50	43.342	32.187	22.027	14.362	8.736	4.078	0
1000		50	50	42.778	30.019	19.344	13.031	8.856	5.184	1.393
1500		50	50	41.521	26.556	17.573	13.37	10.387	7.355	3.891
2000		50	50	37.965	29.813	20.754	16.188	13.318	10.465	7.052
2500		50	50	46.834	33.362	21.963	18.844	16.342	13.904	10.481
3000		50	50	47.556	35.499	27.272	21.643	18.662	17.135	13.567
3500		50	50	41.369	36.953	31.989	25.381	21.795	19.327	15.834
4000		50	50	46.355	40.014	35.531	29.021	24.804	21.693	17.694
4500		50	50	50	44.142	38.539	32.176	27.253	23.548	19.001
5000		50	50	50	49.352	40.657	34.287	28.749	24.085	19.674

Reset to Previous Lookup Table Versions

To reset the normalizer or lookup table to a previous version, select **View > History** to open the History display.

- 1 Highlight the previous version that you want to revert to.
- 2 Lookup tables are independent of normalizers, so if you reset a lookup table to a previous version you must also reset the normalizers to that version (if they have changed). To reset normalizers to match the reverted lookup table version, select the **Reset normalizers** check box.
- 3 Click **Reset**.
- 4 Click **Close** to see the updated lookup table view.

To remove previous versions of the object or comments,

- 1 Highlight the version that you want to remove.
- 2 Click **Remove**.

Click **Clear** to remove all lookup table history.

Adding and Editing Comments About Versions

To add comments,

- 1 Click **Add**.
- 2 In the dialog box, enter your comment.
- 3 Click **OK**. A new History set point is added when you add a comment.

To edit comments,

- 1 Select the comment that you want to edit.
- 2 Click **Edit comment**.
- 3 In the dialog box, edit the comment.
- 4 Click **OK**.

Compare Versions

To compare two different versions of a normalizer or lookup table, highlight the two versions using **Ctrl+click**. Note the following:

- The lower pane shows the difference between the later and the earlier versions.
- Cells that have no entries have no difference.
- Cells that have red entries have a higher value in the later version.
- Cells that have blue entries have a lower value in the earlier version.

History for S_Table

Version	Comment / Action	Date and Time
3	Extrapolated	29-Nov-2013 16:44:46
2	Values filled from optimization output Sum_BTQ_Optimization_Output...	30-Apr-2009 13:26:01
1	Table for 'S'	30-Apr-2009 11:55:38

Reset
Add...
Remove
Edit...
Clear

Reset normalizer

N \ L	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
500			-1.792					-0.366	
1000		-0.306	-2.812					-0.491	-1.36
1500		-0.609	-2.996					-0.716	-1.876
2000			2.341					-1.136	-2.686
2500			-3.648					-1.997	-3.88
3000			-2.346	0.087	-0.788	-0.087	-0.534	-3.429	-5.163
3500			5.782	1.432	-1.457	-0.7	-1.77	-4.093	-5.918
4000			2.961	1.191	-1.688	-1.424	-2.649	-4.817	-6.309
4500				-0.306	-1.855	-1.947	-2.932	-4.91	-6.086
5000				-3.082	-1.444	-1.659	-2.317	-3.637	-5.169

Close Help

See Also

More About

- “Edit Lookup Tables” on page 3-7
- “Lookup Table Properties” on page 3-24
- “Lookup Table Normalizers” on page 3-29


Calibration Manager

In this section...

“Set Up Lookup Tables from a Calibration File” on page 3-20

“Set Up Lookup Tables Manually” on page 3-22

“Copy Lookup Table Data from Other Sources” on page 3-23


Use the Calibration Manager dialog box to manage the sizes, values, and precision of all items that can be calibrated. Open this tool by selecting a lookup table, then selecting **Tools > Calibration Manager** or by clicking the Calibration Manager button  on the toolbar.



You can either set up your lookup tables manually or from a calibration file. You can also copy lookup table data from other sources.

You can enter the required inputs, number of rows and columns and an initial value for lookup table cells when you add a lookup table. Use the **File > New** menu items to make new lookup tables.

Set Up Lookup Tables from a Calibration File


Setting up lookup tables with a calibration file involves two steps.

Step	Description
Import calibration files into the Calibration Manager	<p>You can import calibration files from the CAGE Browser by selecting File > Import > Calibration. Your selected file opens in the Calibration Manager.</p> <p>You can also open the Calibration Manager and import calibration files from within the Calibration Manager window by using the following procedure:</p> <ol style="list-style-type: none"> 1 In the Calibration Manager, open the file by clicking the Open Calibration File button  in the toolbar. <p>The Import Calibration Data dialog box opens.</p> <ol style="list-style-type: none"> 2 Browse to the calibration file, select it, and click Open. <p>In the Calibration Manager, review the files you have imported:</p> <ul style="list-style-type: none"> • Your imported calibration file items appear in the Calibration File Contents pane at the top right. • The lookup tables, normalizers, and other items in your project appear on the left in the Project Calibration Items pane. • The values of the selected item appear in the lower pane, so you can inspect the values in your calibration file and current project lookup tables. Because the import process filters out empty data, any empty variables do not appear.

Step	Description
Import calibration file values into a Lookup Table	<p>To import the data in your calibration file into a lookup table in your project:</p> <ol style="list-style-type: none"> <li data-bbox="628 384 1472 478">1 Click to select both the lookup table in the Calibration File Contents pane and the lookup table in the Project Calibration Items pane with which you want to associate it. <li data-bbox="628 489 1472 636">2 Associate these two items by clicking the button . The values in the calibration file load into the lookup table. You can inspect the values in the lower pane by clicking to select the lookup table in the Project Calibration Items pane. To associate all the items listed in the Project Calibration Items pane with items that have the <i>same names</i> listed in the Calibration File Contents pane, click the Auto button . <li data-bbox="628 972 1472 1003">3 Check the display of your lookup table, and then click Close. <p>To find particular names in a large calibration file, you can click the Calibration File Contents list and type the first few letters of the item that you want to find. The cursor moves to the letters specified as you type.</p>

When you close the Calibration Manager, you can view your updated lookup tables in CAGE. If you want to compare or revert to an earlier version of your lookup table, select **View > History**.

This figure shows the Calibration Manager.




Check the display of your table

Set Up Lookup Tables Manually

To set up the lookup tables manually:


- 1 Select the normalizer or lookup table to set up from the list on the left.
- 2 Enter the number of rows and columns in the edit boxes on the left and select initial values for each cell in the lookup table.
- 3 Click **Apply**.




When initializing lookup tables for a feature calibration, consider your strategy. CAGE cannot fill those lookup tables if you try to divide by zero. Initialize the modifier lookup tables with a value of 1 for all cells if they are multipliers, and a value of 0 if they are added to other lookup tables.


- 4 Check the display of your lookup table, then click **Close**.

Copy Lookup Table Data from Other Sources

You can paste lookup table values from other applications by copying the array in the other application and clicking Paste  in the Calibration Manager:

- 1 Open the desired file and copy the array that you want to import.
- 2 In the Calibration Manager dialog box, click Paste .

You can also set up a lookup table from a text file:

- 1 Click Set Up From ASCII File  in the toolbar.
- 2 Select the desired file, then click **Open**.

If the size of the lookup table is different from the file that you are copying, CAGE changes the size of the lookup table in the session.

See Also

More About

- “Compare and Revert Lookup Tables” on page 3-17
- “Edit Lookup Tables” on page 3-7
- “Initialize Table Values” on page 4-14

Lookup Table Properties

In the Lookup Tables view, to reach the Table Properties dialog box:

- Right-click a table node and select **Properties**.
- Select a lookup table, then select **Table > Properties**

Table Properties: General Tab

The selected lookup table name, type and number of inputs are displayed.

Use the **Table value limits** edit boxes to set a range of values restricting the values in the lookup table.

When you are done, click **OK**.

Table Properties: Table Values Precision Tab

The Table Values Precision tab contains the same settings as the Edit Precision dialog box (reached by clicking the **Edit Precision** button in the Calibration Manager dialog box).

These settings allow you to edit the precision of the number in selected lookup tables and normalizers according to the way lookup tables are implemented in the electronic control unit (ECU). The ECU designer chooses the type of precision for each element to make best use of available memory or processor power.

To edit the precision of a lookup table or normalizer,

- 1 Clear the **Read-only** check box to make the precision writable.
- 2 Select the **Precision type** you require for the lookup table:
 - Floating Point
 - Polynomial Ratio, Fixed Point
 - Lookup Table, Fixed Point

Floating-Point Precision


The advantage of using floating-point precision is the large range of numbers that you can use, but that makes the computation slower.

There are three types of floating-point precision that you can choose from:

- **IEEE double precision** (64 bit)
- **IEEE single precision** (32 bit)
- **Custom precision**

If you choose **Custom precision**, you must specify the following:

- Number of mantissa bits
- Number of exponent bits



Polynomial Ratio, Fixed Point

The advantage of using fixed-point precision is the reduction in computation needed for such numbers. However, it restricts the numbers available to the user.

For example, the polynomial ratio is of the form (see the ratio shown)

$$y = \frac{50x + 0}{0 + 255}$$

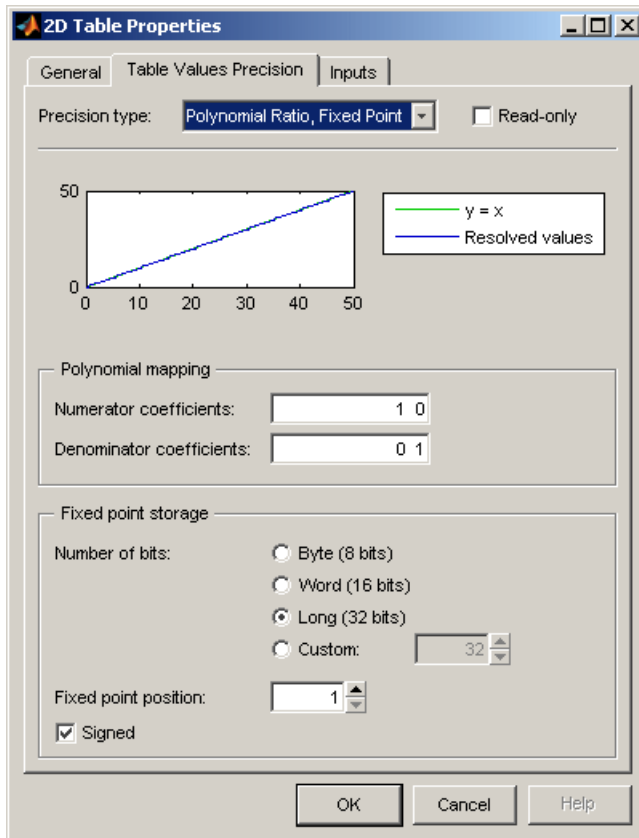
To edit the polynomial ratio,

- 1 Select the **Numerator Coefficients** edit box and enter the coefficients. In the preceding example, enter 500.


The number of coefficients determines the order of the polynomial, and the coefficients are ordered from greatest to least.

- 2 Select the **Denominator Coefficients** edit box and enter the coefficients. In the preceding example, enter 0255.
- 3 To edit the size of the precision, choose from
 - **BYTE** (8 bits)
 - **WORD** (16 bits)
 - **LONG** (32 bits)

- **CUSTOM** (Enter the number of bits in the edit box)
- 4 Select the **Signed** check box if you want the numbers to be negative and positive.



Lookup Table, Fixed Point



The advantage of using fixed-point precision is the reduction in computation needed for such numbers. However, it restricts the numbers available to the user.

For example, consider using a lookup table for the physical quantity *spark advance for maximum brake torque (MBT spark)*. Typically, the range of values of MBT spark is 0–50 degrees. This is the physical data. The ECU can only store bytes of information and you want to restrict the hardware store to a range of 0–8, with at most one decimal place stored.

To adjust the fixed-point precision of the lookup table:

- 1 Select the **Physical Data** edit box and enter the range of the physical data.
- 2 Select the **Hardware Data** and enter the range to store.
- 3 To edit the size of the precision, choose from
 - **BYTE** (8 bits)
 - **WORD** (16 bits)
 - **LONG** (32 bits)
 - **CUSTOM** (Enter the number of bits in the edit box)
- 4 Select the **Signed** check box if you want the numbers to be negative and positive.

Table Properties: Inputs Tab

This tab displays the inputs and variable dependencies for the selected lookup table.

References

[1] Moler, C., Floating Points: IEEE Standard Unifies Arithmetic Model, *The MathWorks Company Newsletter*, 1996.

See Also

More About

- “Edit Lookup Tables” on page 3-7
- “Lookup Table Normalizers” on page 3-29

Lookup Table Normalizers

In this section...
“Editing Breakpoints” on page 3-30
“Input/Output Display” on page 3-31
“Normalizer Display” on page 3-32
“Breakpoint Spacing Display” on page 3-32

What are normalizers? A normalizer is the axis of your lookup table. It is the same as the collection of the breakpoints in your table.

CAGE distinguishes between the normalizers and the lookup tables that they belong to. Using models to calibrate lookup tables enables you to perform analysis of the models to determine where to place the breakpoints in a normalizer. This is a powerful analytical process.

It is important to stress that in CAGE a lookup table can be either one-dimensional or two-dimensional. One-dimensional tables are sometimes known as characteristic lines or functions. Two-dimensional tables are also known as characteristic maps or tables. This is important because normalizers are similar to characteristic lines.


For example, a simple strategy to calibrate the behavior of torque in an engine might have a two-dimensional lookup table in speed and relative air charge (a measure of the load). Also, this strategy might account for the factors of air/fuel ratio (AFR) and spark angle. Each of these compensating factors is accounted for by the use of a simple characteristic line. In CAGE, these characteristic lines are one-dimensional lookup tables. In the example strategy, there are the following lookup tables and normalizers:

- One characteristic map: the torque table
- Six characteristic lines:
 - Two tables: one for AFR and one for spark angle
 - Four normalizer functions: speed, load, AFR, and spark angle

Notice also that a breakpoint is a point on the normalizer where you set values for the lookup table.

Thus, when you *calibrate a normalizer* you place the individual breakpoints over the range of the table axis.

The normalizer node shows the **Normalizer** view, which displays:



Editing Breakpoints

To edit breakpoints:

- Double-click a cell in the **Input** or **Output** column and edit the value.
- Click and drag a breakpoint in the **Normalizer Display** graph or the **Breakpoint Spacing** display.

To view the history of the normalizer function, select **View > History** from the menu. This opens the History dialog box where you can view and revert to previous versions.

Locking and Unlocking Breakpoints

Locking breakpoints ensures that the locked breakpoint does not alter. You might want to lock a breakpoint when you are satisfied that it has the correct value.

To lock a breakpoint, do one of the following:

- Right-click the selected breakpoint in the **Input/Output** display and select **Lock**. Locked breakpoint cells have padlock icons.
- Right-click the selected breakpoint in the **Normalizer Display** or **Breakpoint Spacing** display and select **Lock Breakpoint**. Locked breakpoints are black.

Similarly use the right-click context menus to unlock breakpoints.

Deleting Breakpoints

Deleting breakpoints removes them from the normalizer lookup table. There are still lookup table values for the deleted breakpoints: CAGE determines the positions of the deleted breakpoints by spacing them linearly by interpolation between the non-deleted breakpoints.

Deleting breakpoints frees ECU memory. For example, a speed normalizer runs from 500 rpm to 5500 rpm. Six breakpoints are spaced evenly over the range of speed, that is, at 500, 1500, 2500, 3500, 4500, and 5500 rpm. If you delete all the breakpoints except the endpoints, 500 and 5500 rpm, you reduce the amount stored in the ECU memory. The ECU calculates where to place the breakpoints by linearly spacing the breakpoints between the 500 rpm breakpoint and the 5500 rpm breakpoint.

To delete a breakpoint, right-click the breakpoint and select **Delete Breakpoint**.

Deleted breakpoints are green in the **Breakpoint Spacing** display. You can restore them by right-clicking and selecting **Add Breakpoint**.

Input/Output Display

Input	Output
500	0
1055	1
1609	2
2164	3
2718	4
3273	5
3828	6
4332	7
4836	8
5391	9
5895	10
6500	11

The lookup table consists of the breakpoints of the normalizer function.

The lookup table has inputs and outputs:

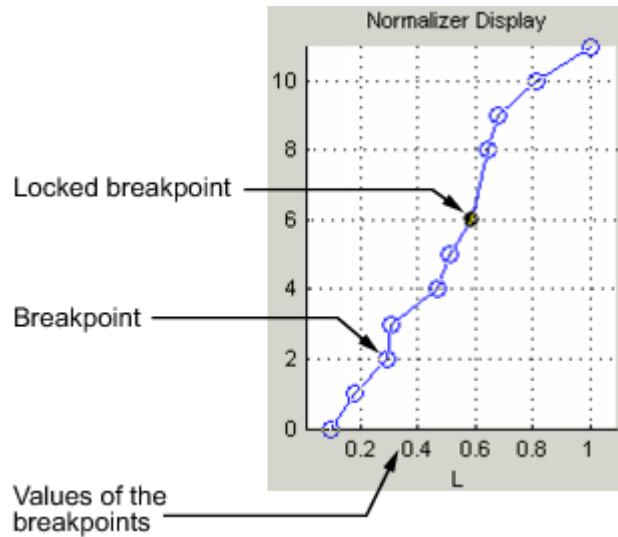
- The inputs are the values of the breakpoints.
- The outputs refer to the row/column indices of the attached lookup table.

To change values of the normalizers in the **Input Output** display, double-click a cell in the **Input** column and change its value.

Normalizer Display

This displays the values of the breakpoints plotted against the marker numbers of the lookup table (that is, the inputs against the outputs).

Click and drag the breakpoints to move them.




Breakpoint Spacing Display

The **Breakpoint Spacing** display shows

- A slice through the model in blue (when feature calibrating regarding a model)
- The breakpoints in red

To move breakpoints, click and drag.



Show Model Curvature

You might want to view the curvature of the model to manually move breakpoints to where the model curvature is greatest.

To display the model slice as its second-order derivative, the curvature of the model,


- Right-click the model in the **Breakpoint Spacing** display and select **Display > Model Curvature**.

You can revert to displaying the model by selecting **Display > Model** from the context menu.

Multiple Slice View


By default the **Breakpoint Spacing** display shows one slice through the model, shown.

Slice Through a Model Surface



Viewing many slices of the model gives a better impression of the curvature of the model. For example, see the following figure.

Many Slices Through a Model Surface



To view multiple slices through the model,

- Right-click the model slice in the **Breakpoint Spacing** display and select **Number of Lines** and choose the number of slices that you want to view from the list.


See Also

More About

- “Edit Lookup Tables” on page 3-7
- “Lookup Table Properties” on page 3-24
- “Compare and Revert Lookup Tables” on page 3-17

Invert Lookup Tables


You can use CAGE to produce a lookup table that is the inverse of another lookup table. Inverting involves swapping a lookup table input with a lookup table output, and you can invert 1-D or 2-D lookup tables.



Inverting a lookup table allows you to link a *forward strategy* to a *backward strategy*; that is, swapping inputs and outputs. This process is desirable when you have a "forward" strategy, for example predicting torque as a function of speed and load, and you want to reverse this relationship in a "backward strategy" to find out what value of load would give a particular torque at a certain speed.

Normally you fill lookup tables in CAGE by comparing with data or models. Ideally you want to fill using the correct strategy, but that might not be possible to find or measure. If you only have a forward strategy but want a backward one, you can fill using the forward strategy (tables or model) and then invert the lookup table.

For example, to fill a lookup table normally from a model, you need the model response to be the lookup table output, and the model inputs to be a function of the lookup table inputs (or it should be possible to derive the input, for example, air mass from manifold pressure). If the available model is "inverted" (the model response is a lookup table input and the lookup table output is a model input) and you cannot change the model, you can invert the lookup table in CAGE.




In the diagram of a lookup table shown, the x- and y-axes represent the normalizers (which you want to be spark and load) and the z-axis is the output at each breakpoint (torque). Filling the lookup table correctly from the model is a two-step process. First, fill a lookup table that has the same input and output as the model, and then fill a second lookup table by inversion.

For the inversion to be deterministic and accurate, the lookup table to be inverted must be monotonic; that is, always increasing or decreasing. This requirement is similar to the one-dimensional example. Every point on the y-axis must correspond to a unique point on the x-axis. The same problem applies also to two-dimensional lookup tables: for any given output in the first lookup

table, there must be a unique input condition; that is, every point on the z-axis should correspond to a unique point in the x-y plane. Some lookup table inversions have multiple values and so do not meet this requirement, just as the square root function can take either positive or negative values. You can use the inversion wizard in CAGE to handle this problem; you can control the inversion process and determine what to do in these cases.

The following example illustrates a lookup table with multiple values. There are two solutions for a single value of torque. CAGE has a lookup table inversion tool that can help overcome this problem. You can specify whether you want to use the upper or lower values for filling certain parts of the lookup table; this allows you to successfully invert a multiple-valued function. See the inversion instructions for 1-D and 2-D lookup tables in the next sections.



The process of inverting a one-dimensional lookup table is different from the process of inverting a two-dimensional lookup table.

Inverting One-Dimensional Lookup Tables

To invert a one-dimensional lookup table,

- 1 Ensure that your session contains two lookup tables:
 - a The first 1D lookup table that you want to invert, filled.
 - b Another 1D lookup table with a different input, which you want to fill with the inverted lookup table.

2 Highlight the lookup table that you want to invert.

3 Click **F¹⁰** or select **Table > Invert**.

The Table Inversion dialog box appears.

4 CAGE selects a suitable table in the **Table to fill** list. Check that this is the lookup table that you want to fill, or select another.

5 The next control, **Resolve non unique inverses with**, specifies what CAGE should do if it encounters multiple values. To minimize the error, leave the default, **Least squares**. The options are:

- **Least squares** selects the range that produces the least error (see below; the last page of the wizard plots the error metric).

- **Minimum** selects the lower of the two if a given number has two possible inverses (like selecting the negative square root of a number).
- **Maximum** selects the uppermost range if a given number has two possible inverses (like selecting the positive square root of a number).
- **Intermediate** selects the middle range if a given number has more than two possible inverses.

For example, the function $y = x^2$ is impossible to invert over the range -1 to 1. You can specify to invert the range from 0 to 1, sacrificing the inversion in the lower range, or the reverse. To select the range from 0 to 1, highlight **Maximum**.

The display shows a comparison between the lookup table (green) and the function $x = f^{-1}(f(x))$.

- 6 The lower pane of the dialog box has a plot that shows the inversion error. If desired, you can change the plot type to view Input and Output. If your forward function is $y = f(x)$, and your inverse function is $x = g(y)$, then, combining these, in an ideal world, you should have $x = g(f(x))$. The plot then displays a line showing x against x and a line showing x against $g(f(x))$. The closeness of these two lines indicates how good the inversion has been: a perfect inverse would show the lines exactly on top of each other.

The plot can show you which part of your lookup table has not successfully inverted and where you should try a different routine.

Note The least squares inversion routine tries to minimize the total distance between these lines. This can sometimes lead to unexpected results. For example, given the function $f(x) = x^2$ from -1 through 1, if you select either positive or negative square root as the inverse, this induces a large error in the combined inverse. If you choose $g(y) = \text{sqrt}(y)$, then $g(f(-1)) = 1$, an error of 2. To minimize this, the least squares routine might choose to send everything to zero and accept a medium error over the whole range rather than a large error over half the range. Use knowledge of the form of the lookup table you are inverting to help you choose which routine to select.

- 7 Click **OK** to accept the inversion or **Cancel** to ignore the result and return to the original lookup table.

Inverting Two-Dimensional Lookup Tables

To invert a two-dimensional lookup table,

- 1 Ensure that your session contains two lookup tables:
 - a The first lookup table you want to invert.
 - b A second lookup table with exactly one of the same inputs, which you want to fill with the inverted lookup table.
- 2 Highlight the lookup table you want to invert.
- 3 Click **F⁴** or select **Table > Invert**.

The Table Inversion dialog box appears.

- 4 CAGE selects a suitable table in the **Table to fill** list. Check that this is the table that you want to fill, or select another.

- 5 The next control, **Resolve non unique inverses with**, specifies what CAGE should do if it encounters multiple values. To minimize the error, leave the default, Least squares. See inverting a 1D table for the other options.
- 6 The lower pane of the dialog box has a plot that shows the inversion error. If desired, you can change the plot type to view the Input and Output. If the forward function is $z = f(x,y)$, and the inverse function is $x = g(y,z)$, then, combining these, in an ideal world you should have $x = g(y,f(x,y))$. The plot then displays a plane showing x plotted against x and y , and a colored surface showing $g(y,f(x,y))$ plotted against x and y . The closeness of these two planes indicates how good the inversion is. If you want to improve the inversion in some areas, try other settings in **Resolve non unique inverses with**.

Click **OK** to accept the result or **Cancel** to ignore the result and return to the original table.

See Also

More About

- “Edit Lookup Tables” on page 3-7
- “Lookup Table Properties” on page 3-24
- “Lookup Table Normalizers” on page 3-29

Import and Export Calibrations

In this section...
“Formats” on page 3-39
“Import Calibrations” on page 3-39
“Export Calibrations” on page 3-39

Formats

You can import and export calibrations in these formats.

- ATI Vision MAT file
- INCA DCM file
- Matrix MAT file
- Simulink Data
- Simple CSV file
- Structure MAT file
- Simple MATLAB file

Import Calibrations

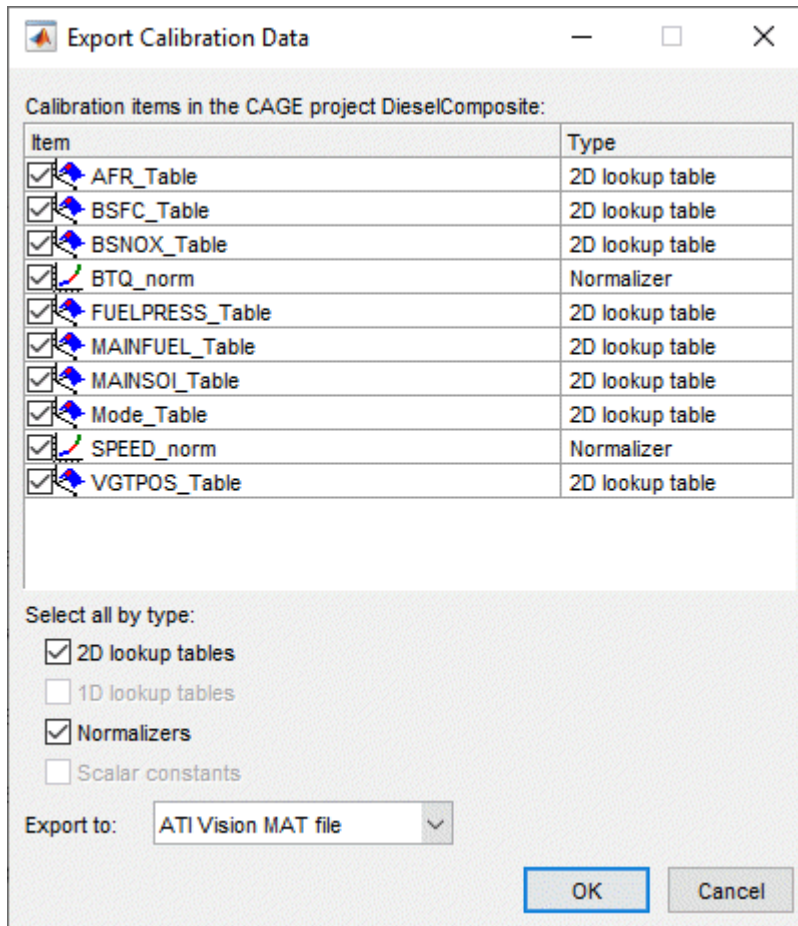
- 1 Select a table.
- 2 Select **File > Import > Calibration** .
- 3 If importing a file, a file browser dialog opens.
 - a Select the type of file you want from the **Files of type** drop-down list, or leave the default **All files (*.*)** and CAGE will try to load the file based on the file extension.
 - b Browse to the file and click **Open** to import.

Export Calibrations

To export your calibration data,

- 1 Select **File > Export > Calibration > Selected Item** or **All Items**, or use the toolbar button to export all calibration items.

The Export Calibration Data dialog appears, for example:



- 2 Select the check boxes of the calibration items you want to export.

You can select all items of a single type by using the check boxes under the list — for example select the **2D lookup tables** check box to select the check boxes of all 2D lookup tables in the list.

If you chose to export **All Items**, all tables, normalizers, curves and constants in the project are available in the list of calibration items.

If you chose to export **Selected Item**, the list items depend on which tree node you have selected. For a table node, the list contains the table and its normalizers. For a Feature or Tradeoff node, the list includes the whole feature or tradeoff (all tables, normalizers, curves and constants). For an optimization node, the list contains any tables filled from the optimization results.

- 3 Use the **Export to** parameter to specify the format. Click **OK**.

- If you select **Simulink Data**, CAGE prompts you to open a Simulink model.
- If you select **Matrix MAT file**, the table data format is compatible with Simulink lookup tables.

See Also

More About

- “Edit Lookup Tables” on page 3-7
- “Calibrate a Feature for an Estimator Subsystem in an ECU” on page 5-2

Feature Calibrations

“Initialize Feature and Calibration Tables” on page 4-12

About Feature Calibrations and Strategies

A feature calibration is the process of calibrating lookup tables by comparing an electronic control unit (ECU) strategy to a statistical model. A Simulink model represents the ECU strategy.

The Model-Based Calibration Toolbox calibrates an estimator, or feature, for a control subsystem in an ECU. Features are usually algebraic collections of one or more lookup tables. You use the features to estimate signals in the engine that are unmeasurable, or expensive to measure, and are important for control. The toolbox can calibrate the ECU subsystem by directly comparing it with a plant model of the same feature.

There are advantages to feature calibration versus calibrating using experimental data. Feature calibrations:

- Smooth noisy data, that is data with measurement error.
- Use statistical models that can make predictions in operating conditions where you have data. You can accurately calibrate and reduce effort gathering experimental data.

A strategy is an algebraic collection of lookup tables, and forms the structure of the feature. Model-Based Calibration Toolbox uses the strategy to estimate signals in the engine that cannot be measured.

For example, a simple strategy to calibrate a feature for Model-Based Calibration Toolbox has two lookup tables:

- Table ranging over the variables speed and load
- Table to account for the behavior of the model as the air-to-fuel ratio (AFR) varies

To evaluate the feature side by side with the model, you need to have a strategy that takes some or all of the same variables as the model. The strategy is expressed using Simulink diagrams. You can either import a strategy or you can construct a strategy.

See Also


More About

- “Process for Calibrating and Filling Feature Tables” on page 4-3
- “Calibrate a Feature for an Estimator Subsystem in an ECU” on page 5-2
- “Fill a Single Lookup Table From a Model” on page 3-14

Process for Calibrating and Filling Feature Tables

A feature calibration is the process of calibrating lookup tables and their normalizers by comparing a collection of lookup tables to a statistical model. The strategy determines the collection of lookup tables. A feature refers to the object that contains the statistical model and the collection of lookup tables.

This is the process for calibrating, filling, and validating feature lookup tables for a subsystem that estimates engine torque.



Step	Description
1	<p>“Set Up Feature and Strategy” on page 4-5</p> <p>Add a feature and construct or import a strategy from a Simulink model file.</p>
2	<p>“Set Up Models” on page 2-11</p> <p>Import and select a response model that relates to the feature that you are filling.</p> <p>A Model-Based Calibration Toolbox response model is a function of a set of variables. Typically, you construct a model using the Model Browser.</p>
3	<p>“Set Up Variables and Constants” on page 2-7 (optional)</p> <p>Optionally, align the feature inputs to a response model. You can use variables, constants, and aliases to create function models.</p>
4	<p>“Initialize Feature and Calibration Tables” on page 4-12 (optional)</p> <p>Optionally, initialize the feature model lookup table parameters.</p>

Step		Description
5	"Fill Feature Lookup Tables" on page 4-18	Initialize and optimize your feature model lookup table parameters.
6	"Validate Feature Fill" on page 4-28	Validate the feature fill by comparing the feature model and response model outputs at the same input values.
7	"Import and Export Calibrations" on page 3-39	Export feature calibration tables to a file.

See Also

More About

- "About Feature Calibrations and Strategies" on page 4-2
- "Calibrate a Feature for an Estimator Subsystem in an ECU" on page 5-2
- "Fill a Single Lookup Table From a Model" on page 3-14

Set Up Feature and Strategy

A feature calibration is the process of calibrating lookup tables by comparing an electronic control unit (ECU) strategy to a statistical model. A Simulink model represents the ECU strategy. For more information about the process, see “Process for Calibrating and Filling Feature Tables” on page 4-3.

The Model-Based Calibration Toolbox calibrates an estimator, or feature, for a control subsystem in an ECU. Features are usually algebraic collections of one or more lookup tables. You use the features to estimate signals in the engine that are unmeasurable, or expensive to measure, and are important for control. The toolbox can calibrate the ECU subsystem by directly comparing it with a plant model of the same feature.

A strategy is an algebraic collection of lookup tables, and forms the structure of the feature. Model-Based Calibration Toolbox uses the strategy to estimate signals in the engine that cannot be measured.



To evaluate the feature side by side with the model, you need to have a strategy that takes some or all of the same variables as the model. The strategy is expressed using Simulink diagrams. You can either import a strategy or you can construct a strategy.

You can import features for selected Simulink subsystems.

Add a Feature

A feature consists of a model and a collection of lookup tables organized in a strategy.

To add a feature, select **File > New > Feature**. This automatically switches you to the **Feature Filling** view and adds an empty feature to your session.

An incomplete feature is a feature that does not contain both an assigned model and a strategy. If a feature is incomplete, CAGE displays  in the tree display. If a feature is complete, CAGE displays  in the tree display.

Import a Strategy from Simulink

- 1 Highlight the top feature node in the tree display.
- 2 Select **File > Import > Strategy**.
- 3 Select the appropriate Simulink model file.

If there is a single outer output, CAGE automatically imports the strategy.

- 4 If there are multiple outer outputs, CAGE prompts you to import all output ports into separate features.
- 5 Use the MBC Feature Importer to preview lookup tables and import features for selected Simulink subsystems. Click **OK**.
- 6 To view a text representation of your strategy, select the **Feature** node.

If parsing the Simulink diagram fails and you see an error message, then the toolbox discards changes in the current CAGE project. Correct the Simulink model and reparse it.

Model Structure and Strategy Hierarchy

CAGE uses the Subsystem hierarchy in the Simulink model to generate subfeatures in CAGE. This makes it easier to understand the structure of the strategy and relate it to the Simulink model. However, be aware how CAGE creates subfeatures from Simulink models:

- When a subsystem has more than one output, CAGE adds the output name to the subsystem for the subfeature name. Rename an output before importing if you want a particular name for the CAGE subfeature. CAGE creates unique names.
- CAGE creates a subfeature from outputs in subsystems. CAGE works backward from outputs and includes all input blocks to the output in the subfeature, including blocks outside the subsystem. Subfeatures are not identical to Simulink subsystems.
- CAGE shows the subfeature hierarchy in the Feature tree. Each subfeature is also visible at the top level of the tree.

Lookup Tables, Normalizers, and Constants

Names and reuse	<ul style="list-style-type: none"> • If the parameter is a valid variable name and not a MATLAB expression, CAGE uses the variable names to determine the lookup table and normalizer names. This supports lookup table reuse and avoids the need for explicit normalizer blocks. • For constants, CAGE uses the block name or constant variable name, if defined. Constants are reused if they have the same name as an existing constant and the value is the same. • If a lookup table with the same name exists in the project and the input expressions for these lookup tables are the same, then the toolbox reuses the lookup table. Similarly, the toolbox reuses normalizers if they have the same inputs as the existing normalizer. <p>If a lookup table of the same name exists in project and the new lookup table has different inputs, then CAGE asks what option you want:</p> <ul style="list-style-type: none"> • Create a lookup table with a different name (suffix <code>_1</code>). • Reconnect the lookup table inputs using the current Simulink block connections. • Cancel. You can then edit the model to resolve differences if desired.
-----------------	---

Lookup table and normalizer structure	<ul style="list-style-type: none"> • Shared normalizers can be used as inputs to multiple lookup tables. You can view shared normalizers at the top of the CAGE Lookup Table tree. • You can create 1-D tables with or without normalizers. If you add a CAGE Function block from <code>cgeqlib</code> without a normalizer, a 1-D lookup table with an internal normalizer is created on parsing. If the Function block has a normalizer as its input, then you can use shared normalizers (from the list of available normalizers in CAGE). After creation, you cannot change from using shared normalizers to internal normalizers. • 2-D lookup tables always have shared normalizers. If the input to the Table block is not a normalizer, then CAGE creates a normalizer when the strategy is parsed. <p>CAGE assigns normalizer names using the lookup table breakpoint (or row/column index) variable names, if available. If an expression defines the breakpoints and not a variable, then CAGE names normalizers using the form <code>tablename_normY</code> or <code>tablename_normX</code>.</p> <ul style="list-style-type: none"> • If you change the name of inports, table, or normalizer blocks for blocks associated with existing CAGE items, then the CAGE items name is changed. The name is unique for the current CAGE project (suffix <code>_1</code> is added if necessary to create a unique name). • Prelookup tables must feed into an Interpolation block using a Prelookup block. • Normalizer blocks, if used, must be inputs to lookup tables.
Data import	<ul style="list-style-type: none"> • Table and normalizer data is imported from Simulink. <p>You must be able to run Update Diagram on the Simulink model and the data must have a single source. Otherwise, lookup table data is left empty and you must set up the lookup tables, normalizers, and constants using the Calibration Manager.</p> <ul style="list-style-type: none"> • Constant data is read from Constant or Gain blocks.

Block Support

Blocks	Description
Math Operations	CAGE supports a subset of Simulink blocks. To view the supported blocks, open the <code>cgeqlib</code> library.
Switch blocks	CAGE can import the Simulink Switch block and MultiportSwitch block. When you import the Switch block, CAGE converts it to a CAGE block called IfExpr. You can view the block in the <code>cgeqlib</code> library.
Logical Operator and Relational Operator	Logic and Boolean expressions: CAGE can import the Logical Operator and Relational Operator blocks.
Product and Sum	Polynomial expressions: CAGE builds polynomial expressions using Horner's form (Product and Sum blocks).
Dot Product	Allows the implementation of weighted sum expressions.

Blocks	Description
Interpreted MATLAB Function	Use to implement more general functions in CAGE features. The function must be vectorized, that is, it must accept matrix inputs of the form ([u1 , u2 , u3 , ... , un]).
Signal conditioning	CAGE ignores signal conditioning blocks. Several standard Simulink blocks are intended for conditioning signals, but you can ignore the blocks for the purposes of steady-state analysis in CAGE. These blocks include the Signal Conversion, Rate Transition, Data Type Conversion, and Initial Condition blocks. CAGE ignores them, making it easier to import existing strategy diagrams.
Merge and If	CAGE supports Merge blocks fed by enabled or If/Elseif/Else subsystems, common in industry models.

Only scalar inputs are supported except for the MATLAB Function, Dot Product, Polynomial, and Interpreted MATLAB Function blocks, which accept multiple inputs as inputs to the expression (e.g., $u(1)+u(2)$).

Loop Handling

CAGE cannot handle expressions with loops. If CAGE detects a loop, then CAGE asks if you want to break the loop by introducing a variable called `previousOutputName`. If possible, CAGE tries to break loops at a feature boundary (for example, a top level or subsystem output).

If you decide not to break the loop, an error message informs you which blocks are involved in the loop.

CAGE ignores Delay blocks and Resettable Delay blocks to facilitate loop parsing.

Importing Older Strategies

To parse strategies from previous releases, you can use the function `cgStrategyTables` to set the style for strategy parser behavior. Use `cgStrategyTables` to import any older strategies saved as Simulink model files.

- For strategies from R2013a or later, CAGE interprets the first input to a 2-D lookup table as Y (rows). Previously, CAGE parsed the first input as columns and the second as rows. Use the `cgStrategyTables` backward compatibility modes for previously saved strategies, if needed.
- For strategies from R2008a to R2012b, lookup table blocks are always interpreted as lookup tables, because there are separate lookup and normalizer blocks. In R2008a, a normalizer block was added to the `cgeqlib` library.

To set the parser style to R2008a, use:

```
cgStrategyTables('R2008a')
```

- For strategies older than R2008a, CAGE interprets 1-D lookup tables as normalizers if they feed into a lookup table.

To parse pre-R2008a strategies, use this function:

```
cgStrategyTables('Pre-R2008a')
```

This reverts the parser behavior to the pre-R2008a interpretation of 1-D lookup tables. CAGE issues a warning when converting a 1-D lookup table to a normalizer. You can turn off the warning:

warning off mbc:cgslparser:ObsoleteNormalizer

- To reset the parser style to R2013a and later, use:

```
cgStrategyTables('R2013a')
```

- To query the current strategy table style, enter:


```
Style = cgStrategyTables
```

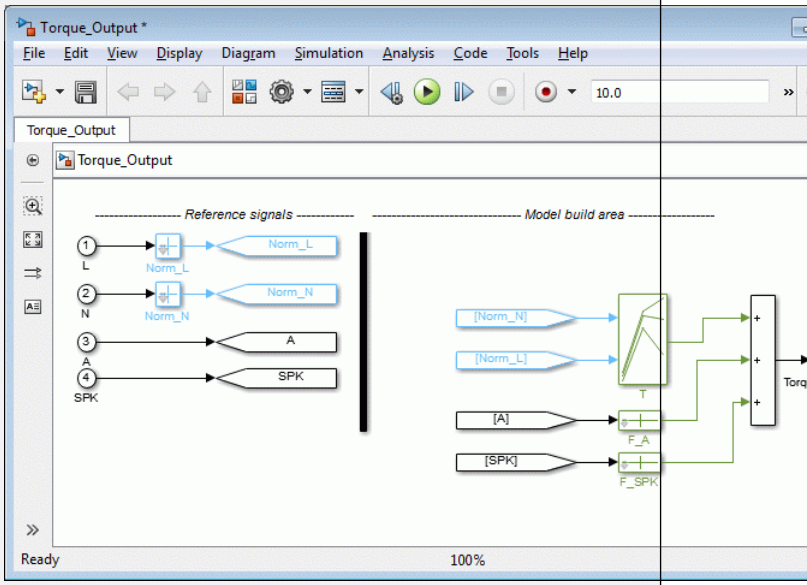
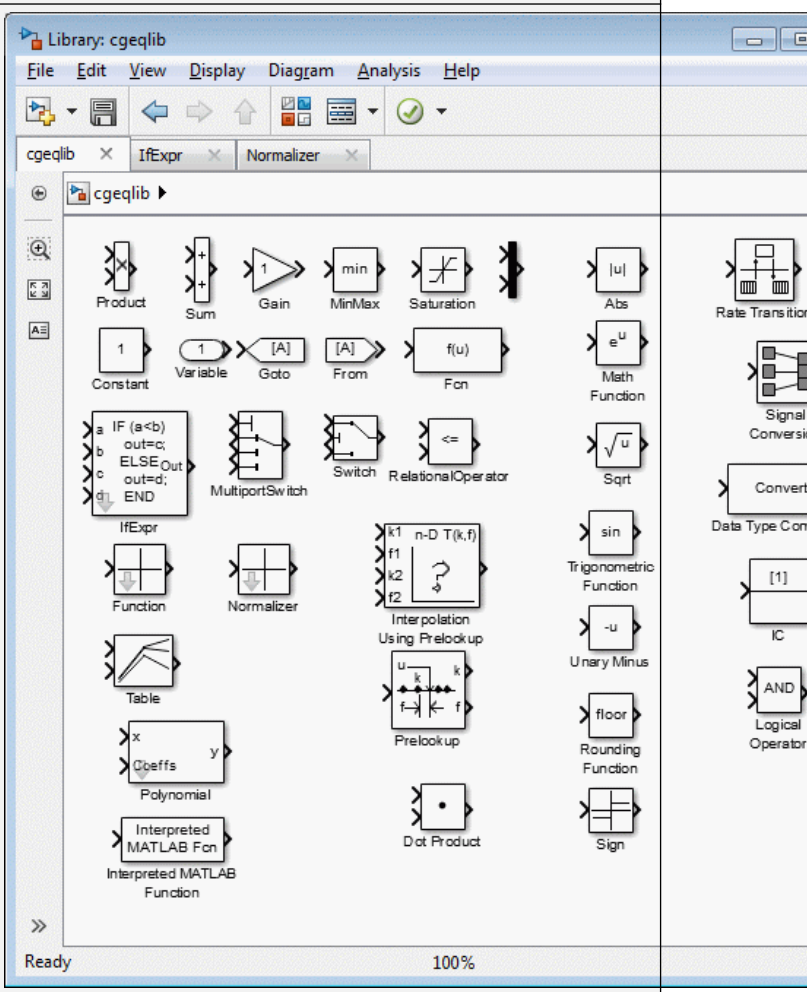
Construct a Strategy

To construct a strategy from CAGE rather than import an existing model:

- 1 Highlight the correct feature by selecting the **Feature** node.
- 2 Select **Feature > Edit Strategy**.

Three Simulink windows open:



Window	Description	
Library, cgeqlib	Blocks available for building a strategy.	 <p>The screenshot shows the 'Torque_Output' window. On the left, under 'Reference signals', there are four inputs: 1 (L), 2 (N), 3 (A), and 4 (SPK). These connect to normalization blocks 'Norm_L' and 'Norm_N'. The 'Model build area' contains a table block 'T' which takes '[Norm_N]' and '[Norm_L]' as inputs. Below it are function blocks 'F_A' and 'F_SPK' which take '[A]' and '[SPK]' as inputs. The outputs of 'T', 'F_A', and 'F_SPK' are summed to produce the final 'Torque' output.</p>
Library	Existing blocks in your CAGE project, organized into libraries. Blocks are From blocks referencing the CAGE items such as tables and normalizers.	 <p>The screenshot shows the 'Library: cgeqlib' window. It displays a grid of various blocks categorized into several groups: <ul style="list-style-type: none"> Basic operations: Product, Sum, Gain, MinMax, Saturation, Abs, Constant, Variable, Goto, From, Fcn. Control and logic: IfExpr, MultipointSwitch, Switch, RelationalOperator, Function, Normalizer. Mathematical functions: Math Function, Sqrt, Trigonometric Function, Unary Minus, Rounding Function, Sign, Interpolation Using Prelookup, Prelookup, Dot Product. Other: Table, Polynomial, Interpreted MATLAB Fcn, Rate Transition, Signal Conversion, Convert, Data Type Conversion, IC, Logical Operator. </p>

- 3 In the strategy window, build your strategy using the blocks in the library windows. To perform a feature calibration, the strategy and the model must have some variables in common.
- 4 To parse the strategy into the CAGE session, double-click the blue output circle.

The `cgeqlib` library contains all the blocks available for building a strategy.

Export Strategies

Export Simulink strategies if you want to:

- Include a strategy in a Simulink vehicle model.
- Compile the strategy using Simulink Coder™ to produce C code.
- Evaluate the strategy using Simulink.

To export a strategy from CAGE:

- 1 Highlight the **Feature** node that contains the strategy that you want to save.
- 2 Select **File > Export > Strategy**.
- 3 Assign a name for your strategy.

The strategy is saved as a Simulink model file.

On export, table data is stored in variables. Indices are written to Simulink parameters using colon expressions.

- For 2-D lookup tables, `0:size (Table,1)-1, 0:size (Table,2)-1, TableName`
- For 1-D lookup tables with shared normalizers, `0:length (Table)-1, TableName`
- For 1-D lookup tables, `NormalizerName` and `TableName`
- `NormalizerName, 0:length(Normalizer)-1`

The data is stored in the model workspace. The model workspace data is copied to the new model when you copy the strategy block. You must be able to run **Update Diagram** on the Simulink model or the copy process does not work. Copying the strategy model requires a Model-Based Calibration Toolbox.

See Also

More About

- “About Feature Calibrations and Strategies” on page 4-2
- “Calibrate a Feature for an Estimator Subsystem in an ECU” on page 5-2
- “Fill Feature Lookup Tables” on page 4-18
- “Initialize Feature and Calibration Tables” on page 4-12

Initialize Feature and Calibration Tables

If you import your strategy with lookup tables already initialized, you do not need to initialize your lookup tables. Use the Feature Fill Wizard to fill your lookup tables. For more information, see “Fill Feature Lookup Tables” on page 4-18.

You can use the Calibration Manager to import lookup table values from external sources and set the size of lookup tables and normalizers. For more information, see “Calibration Manager” on page 3-20.

Initialize a Feature


You can initialize a feature to set the values of the normalizers over the variable range and put specified values into each lookup table cell. A lookup table that is already initialized provides a useful starting point for a more detailed calibration.


For example, a simple feature for maximum brake torque (MBT) consists of the following lookup tables:

- Speed (N), load (L) lookup table
- A lookup table to account for the behavior of air/fuel ratio (A)

Initializing this feature sets the values of the normalizers for speed, load, and AFR over the variable ranges and puts specified values into each lookup table cell.

To initialize the feature, perform these steps:

- 1 Click  to open the Feature Initialization Options dialog box.




- 2 Enter the ranges for the breakpoints in your normalizers.
 - L has range 0.2 to 0.811.
 - N has range 750 to 6500.
 - A has range 11 to 17.6.
- 3 Enter the initial lookup table value for each cell in each lookup table. Above, the cell values are
 - Table T has initial value 0.
 - Table F_A has initial value 0.
- 4 Click **OK** to initialize the feature.

Initialize Breakpoints

You can initialize normalizer breakpoints individually, or initialize the whole feature. Initialization places the breakpoints at even intervals along the range of the variable defined for the normalizer. When you add a lookup table and specify the inputs in the Table Setup dialog box, CAGE initializes the table normalizers by spacing the breakpoints evenly over the input variable ranges. If you have edited breakpoints, you can return to even spacing by using the Initialize function.

To space the breakpoints evenly:


- 1 Click  on the toolbar or select **Normalizer > Initialize**.
- 2 In the dialog box, enter the range of values for the normalizer.
- 3 Click **OK**.

For example, for a torque table with two normalizers of engine speed and load, you can evenly space the normalizer breakpoints over the range between 500 and 6500 rpm for speed and between 0.1 and 1 for relative air charge. In the dialog box, enter 500 6500 for the speed normalizer, N, and 0.1 1 for the load normalizer, L.

Initialize Table Values

You can initialize tables individually or initialize the whole feature. Initializing table values sets the value of every cell in the selected table to a constant.

To initialize the table values:


- 1 Click  or select **Table > Initialize**.
- 2 Select the constant value that you want to insert into each cell.

Considerations:

- Filling with zeros can cause a problem for some strategies that use modifier tables. For example, your strategy might use several speed-load tables for different values of AFR. The strategy might use an AFR table as a modifier to add to a single speed-load table to adjust for the effects of different AFR levels on your torque output.
- If the strategy uses modifier tables that are multipliers, avoid initializing the modifier tables with zeros. Initialization can result in division by zero failures.
- If the table is a modifier that is added to other tables, initially fill it with zeros.
- If the table is a modifier that multiplies other tables, fill it with ones.

Optimize Normalizer Breakpoints

You can use CAGE to space the breakpoints over the range of normalizers. Select a normalizer in the tree display.



To space the breakpoints, either click the buttons on the toolbar or select from the following options on the **Normalizer** menu.

Normalizer	Description
Initialize	Spaces the breakpoints evenly along the normalizer.

Normalizer	Description
Fill	Spaces the breakpoints by reference to the model. For example, you can place most breakpoints where the model curvature is the greatest.
Optimize	Moves the breakpoints to minimize the least square error over the range of the axis. To optimize normalizers, each normalizer must have a single variable input that is an input to the model. For 2-D tables, the single variable input must be different from the other normalizer input. Only available when you are performing Feature calibrations.

Optimize Breakpoints

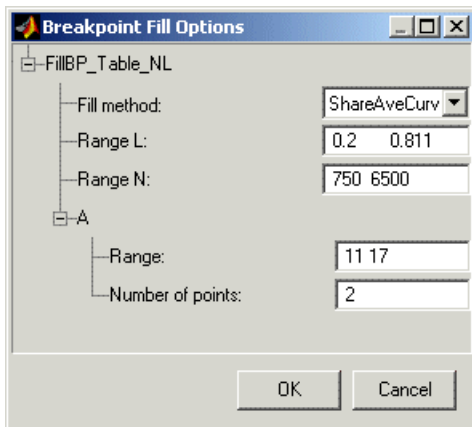
Optimizing breakpoints spaces the breakpoints by reference to the model.

For example, suppose a model of the spark angle that produces the maximum brake torque (MBT) has these inputs: engine speed N , relative air charge L , and air-to-fuel ratio A . You can space the breakpoints for engine speed and relative air charge over the range of these variables by referring to the model.

To space the breakpoints based on model curvature, complete these steps:

- 1 Click  or select **Normalizer > Fill**.

The Breakpoint Fill Options dialog box opens.



- 2 Choose the appropriate method to space your breakpoints, from the drop-down menu in the dialog box.
- 3 Enter the ranges of the values for the normalizers.

The preceding example shows **Range N** 500 6500, and **Range L**, 0.1 1.

- 4 Enter the ranges of the other model variables.

CAGE spaces the breakpoints by reference to the model at selected points. The example shows 11 17 for the **Range** of **A** and 2 for the **Number of points**. This takes two slices through the model at $A = 11$ and $A = 17$. Each slice is a surface in N and L . That is, $MBT(N, L, 11)$ and $MBT(N, L, 17)$.

CAGE computes the average value of these two surfaces to give an average model $MBT_{AV}(N, L)$.

If you set **Number of points** to one, and specify a range, then the mean of the range is chosen as the evaluation point.

5 Click **OK**.

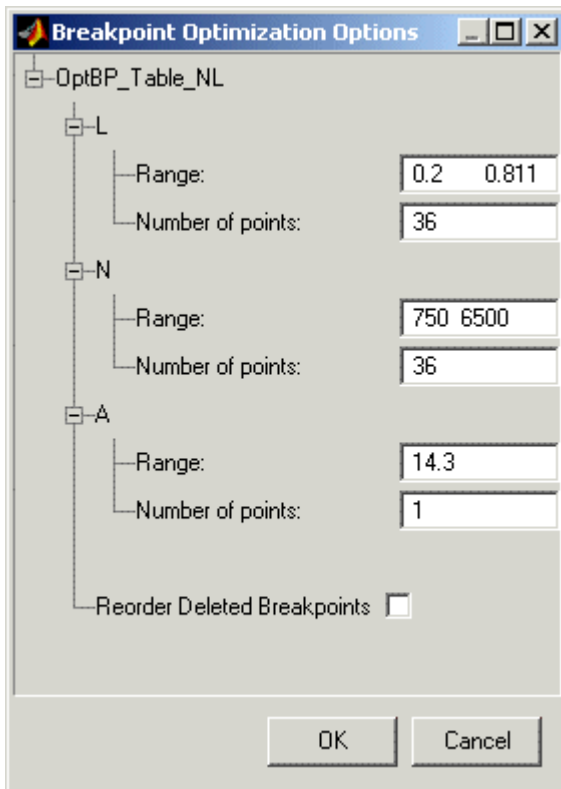
If any of the breakpoints is locked, each group of unlocked breakpoints is independently spaced according to the selected algorithm.

If you increase the number of slices through the model, you increase the computing time required to calculate where to place the breakpoints.

To optimize the breakpoints, complete these steps:

- 1 Ensure that the optimization routine works over reasonable values for the table by choosing one of these methods:
 - a Select **Normalizer > Initialize**.
 - b Select **Normalizer > Fill**.
- 2 Click on the toolbar or select **Normalizer > Optimize**.

This opens the following dialog box.



3 Enter the ranges for the normalizers.

The example shows 0.2 0.811 for the **Range** of **L**, and 750 6500 for **N**.

4 Enter the appropriate number of grid points for the optimization.

This defines a grid over which the optimization works. In the preceding example, the number of grid points is 36 for both L and N . This mesh is combined using cubic splines to approximate the model.

- 5 Enter ranges and numbers of points for the other model variables.

The example shows a **Range of A** of 14.3 and the **Number of points** is 1.

- 6 Decide whether to reorder deleted breakpoints by clicking the option button.

If you choose to reorder deleted breakpoints, the optimization process might redistribute them between other undeleted breakpoints.

- 7 Click **OK**.

CAGE calculates the table filled with the mesh at the current breakpoints. Then CAGE calculates the total square error between the table values and the mesh model.

The breakpoints are adjusted until this error is minimized, using nonlinear least squares optimization.

When optimizing the breakpoints, the default:

- Range for the normalizer variable is the range of the variable.
- Value for all other model variables is the set point of the variable.
- Number of grid points is three times the number of breakpoints.

See Also

fmincon | lsqnonlin

More About

- “Calibration Manager” on page 3-20
- “Calibrate a Feature for an Estimator Subsystem in an ECU” on page 5-2
- “Add, Duplicate, and Delete Tables” on page 3-8
- “Fill Feature Lookup Tables” on page 4-18
- “Set Up Feature and Strategy” on page 4-5
- “Lookup Table Normalizers” on page 3-29

Fill Feature Lookup Tables

Use the Feature Fill Wizard to fill and optimize the values in tables by reference to the model or data. You can fill multiple tables at once using the wizard, and you can **Fill** from the top feature node or from any table node in a feature. Use **Fill** at the top feature node to calibrate the entire feature, that is, fill all the table values by referring to a model.

Before using the Feature Fill wizard:


- Your project must contain a feature and a model or a data set.
- Initialize your tables, unless you imported your strategy with tables already initialized.

Filling and Optimizing Lookup Table Values

When filling a feature table, the Feature Fill Wizard optimizes the table values to minimize the total square error between the feature values and the model. This routine optimizes the fit between your strategy and your model. Using **Fill** places values into your table. The optimization process shifts the cell values up and down to minimize the overall error between the interpolation between the model and the strategy.

To fill feature tables:

- 1 Select **Feature > Fill Feature** to open the Feature Fill Wizard.



Select tables to fill. Select the check boxes of the tables you want to fill. For each table, you can set these options.

Option	Description
Clear Mask	Select to clear any table mask. Model-Based Calibration Toolbox adds filled table cells to the table mask and uses the entire mas for extrapolation.
Extrapolate	Select to extrapolate across the whole table after filling cells. The extrapolation is based on the filled cells in the mask and any locked cells.
Table Bounds	Set bounds on the table values.

Option	Description
Row Gradient Bounds	Set bounds on the gradient (slope) between rows and columns. For example, entering 0 Inf in the left edit box imposes the constraint that the gradient must be positive (increasing) between successive rows.
Column Gradient Bounds	

2 Set up the feature fill.

Feature Fill Wizard

Set up Fill
Select the item to fill the lookup tables from. Optionally, select a data set or a grid and smoothing factors.

Reference
Feature: Torque_Feature(N, L, A, delta_S)
Filled by: BTQ(S, N, L, ICP, ECP, CylinderMode) Model
 Use percentage error

Input data source
Input data: Grid
Constraint: Boundary constraint of BTQ
Setup... Remove

Smoothing
Smoothing factor: 0
Smoothing for each normalizer:

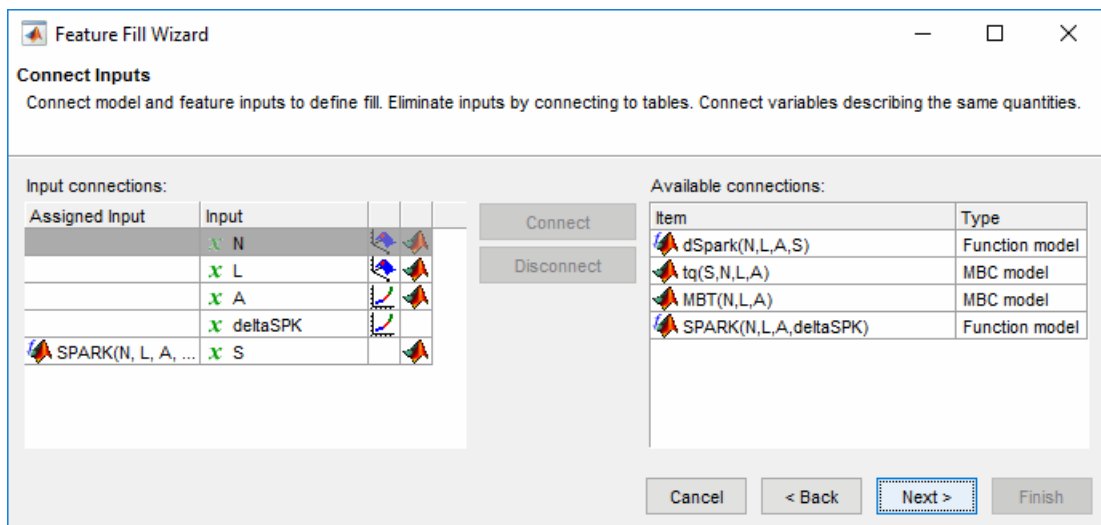
	Row (Y)	Column (X)
F_A	on	N/A
F_SPK	on	N/A
T	on	on

Cancel < Back Next > Finish

Parameter	Description
Reference	Specify the model or data set column to use to fill the feature. By default, the feature fill uses the absolute error. To use the percentage error, select Use percentage error .
Input data source	For Input data , specify either Grid or Data Set. The data set option is available only if there is a data set containing some of the inputs to the feature. For Constraint , you can click Setup to create a constraint. Constraints restrict the input grid or data set values that CAGE uses to fill the feature. By default, the toolbox sets the constraint to the boundary model associated with the main model.


Parameter	Description
Smoothing	<p>Specify a Smoothing factor on table axes.</p> <p>If you are filling the feature with a data set, specify a smoothing factor.</p> <p>The smoothing factor applies to all the tables CAGE fill.</p> <p>The smoothing factor applies a smoothing penalty to the optimization. The smoothness penalty uses the second derivative to avoid steep jumps between adjacent table values.</p>

- 3 Connect the model and feature inputs to define the feature fill.



Click **Next**.

- 4 Set the grid values that you want to use for the feature fill.



Optionally, you can click **Initialize From Normalizer** to use the normalizer breakpoints as a variable value. In the Select Normalizer dialog box, you can also choose to interleave values between breakpoints. Interleaving values can minimize interpolation error by adding values between each normalizer value. In this way, you can create a grid of more points than table cells to optimize over. Select normalizers in the dialog box to use those breakpoints as the value.

- 5 Click **Finish** to fill the tables.

CAGE evaluates the model over the number of grid points specified, then calculates the total square error between this mesh model and the feature values. If no gradient constraints exist, CAGE uses `lsqnonlin` to adjust the table values until the error is minimized. If there are gradient constraints, to specify the gradient of the table at each cell, CAGE uses `fmincon`.

Saving and Reusing Feature Fill Settings

After feature filling, the Feature Fill Wizard saves your settings. You can run and manage your saved fill settings from the Feature Fill Settings pane. If you import new models, CAGE reuses the saved settings for the feature fill.

To	Select
Define new fill settings starting from the defaults	Select Feature > New Fill Setting .
Rerun all saved fill settings in your feature	Select Feature > Run All Fill Settings .
Copy and modify some saved settings	Select Feature > Duplicate Selected Fill Setting .
Delete all saved fill settings for the selected feature	Select Feature > Clear All Fill Settings .

See Also

`fmincon` | `lsqnonlin`

More About

- “Calibrate a Feature for an Estimator Subsystem in an ECU” on page 5-2
- “Fill Tables by Extrapolation” on page 3-9
- “Initialize Feature and Calibration Tables” on page 4-12
- “Validate Feature Fill” on page 4-28

Optimize Calibrations with Transient Controllers

Use the transient feature fill to optimize calibrations with transient controllers, enabling calibration of dynamic systems like diesel particulate filter (DPF) and nitric oxide and nitrogen dioxide (NOx) emission controllers and on-board diagnostics (OBD) estimators. Specifically, CAGE can optimize calibrations for transients features that:

- Include delays, discrete integrators and filters, and feedback loops.
- Contain many (> 50) lookup tables that require estimation.
- Have multiple test data files to tune the feature.
- Require analytic derivatives to perform large-scale optimizations.

For large feature fills, the fill can take a long time to run. If it is available, CAGE uses the Parallel Computing Toolbox™ to evaluate calibration data files on different workers.


To use the transient feature fill, follow these workflow steps.

Step	Description
1	“Import Simulink Model” on page 4-22
2	“Run Optimization” on page 4-23
3	“Restart Optimization” on page 4-25
4	“Analyze Results” on page 4-25
5	“Export Calibration” on page 4-26

Import Simulink Model

Use CAGE to import the Simulink model that has the transient features.

- 1 In CAGE, select **File > Import > Strategy**. Use the browser to select the Simulink model with the transient features.
- 2 In the MBC Feature Importer dialog box, on the **Feature Details** tab, select **Allow transient features**.



Note If you select **Allow transient features**, your model cannot have algebraic loops. Transient feature calibration does allow models with feedback loops. Each loop must have a block with no direct feedthrough.

CAGE supports these transient Simulink blocks:

- Unit Delay
- Delay - Discrete, fixed, and variable
- Discrete-Time Integrator - Forward, backward, and trapezoidal
- Discrete Transfer Fcn - Initial conditions set to zero with no ability to modify the transfer function
- Discrete Filter

Run Optimization

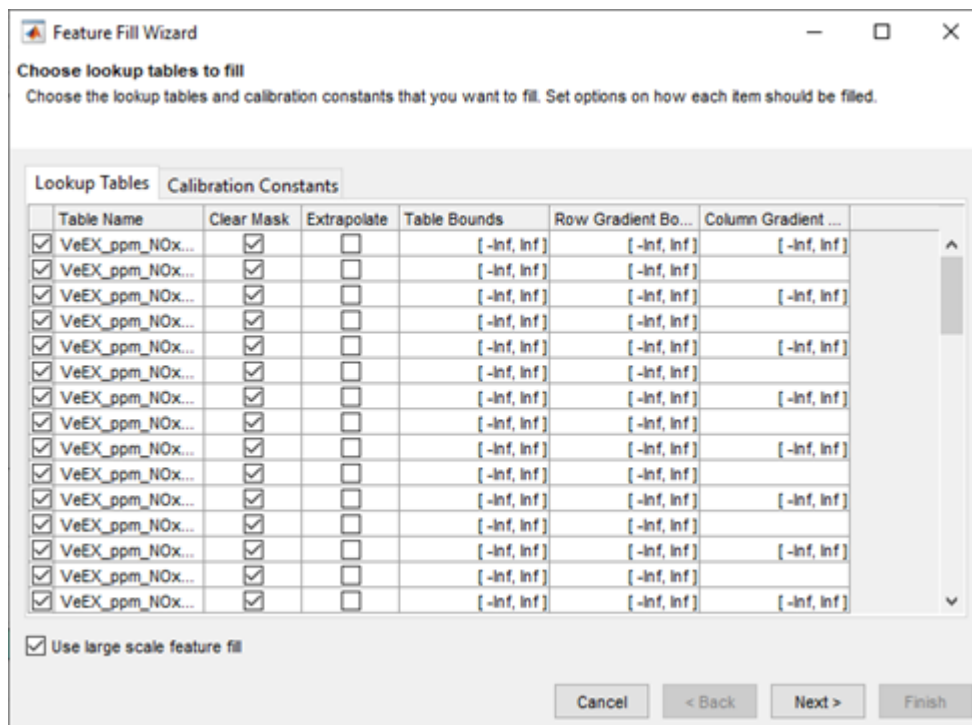
Before you start, make sure that you have the statistical response model aligned with the feature inputs. For more information, see “Set Up Models” on page 2-11 and “Set Up Variables and Constants” on page 2-7.

To run the optimization, in CAGE, click Feature Filling



. Use the Feature Fill Wizard to specify the fill settings.

- 1 Select **Feature > Fill Feature**.
- 2 In the Feature Fill Wizard, select **Use large scale feature fill**.

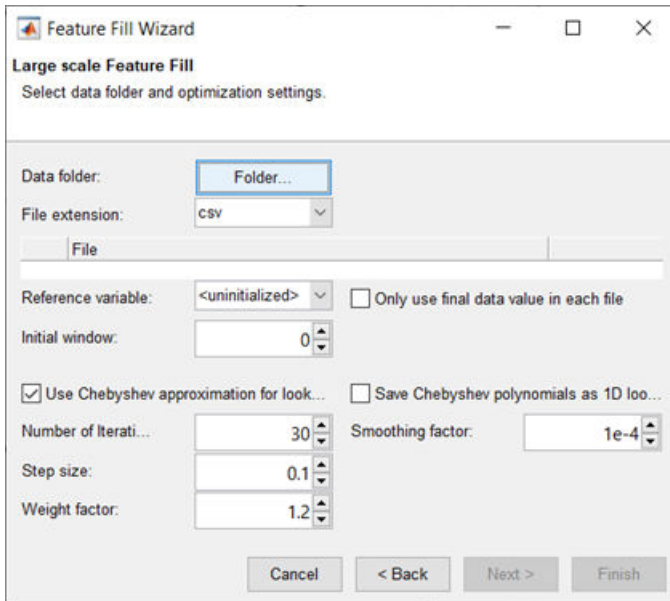


- On the **Lookup Tables** tab, select all of the tables.

- On the **Calibration Constraints** tab, select the calibration constraints that you want to use in the optimization. CAGE does not support bounds or gradient constraints for large-scale optimizations.

Click **Next**.

- 3 In the wizard, specify the **Large scale Feature Fill** optimization settings.



Parameter	Description
Data folder File extension	Folder containing test data and the file extension. The extension is either CSV or xlsx. After you specify the folder and extension, CAGE finds compatible files. To be compatible, each file must contain all the feature inputs and a common additional column for the software to use to match against the feature.
Only use final data file in each file	Select to use the final data for cases that have measured output data only at the end of the test.
Use Chebyshev approximation for lookup tables	Optimize lookup tables by using Chebyshev polynomials rather than directly by using lookup table values. Chebyshev polynomials represent the lookup tables being optimized with fewer parameters than the lookup table cells. This type of optimization is useful when filling many (and possibly large) lookup tables. You can use Chebyshev polynomials to start the filling process and switch to filling directly with lookup table values for the last stage.
Save Chebyshev polynomial as 1D lookup tables	Save Chebyshev polynomial coefficients as 1D lookup tables for visual inspections and export to CSV files.

Parameter	Description
Smoothing factor	CAGE applies a smoothing factor to the tables in the feature fill during the optimization. The smoothness penalty uses the second derivative to avoid steep jumps between adjacent table values.
Weight factor	Weighting factor for each test file. CAGE provides tests with a higher than average error a higher weight. In the case of continuous matching, the weighting factor is inversely proportional to the number of observations in the test.

- 4 Click **Finish**. As the feature fill runs, view progress in the Fill Progress dialog box.
- The top plot displays fit for each file in the current iteration compared to previous iteration.
 - The bottom plot shows the improvement in RMSE over each iteration.
 - View lookup tables in CAGE.

Before accepting an optimization, use the Fill Progress dialog box to review the results. Optimization runs can take a long time, possibly overnight. Click:

- **Accept** to use the lookup table values associated with the best solution found to date.
 - Click **Cancel** if you want to discard the optimization run (for example, if the optimization is unstable). The table values are not changed. If necessary, you can “Restart Optimization” on page 4-25.
- 5 If CAGE displays a warning message about breakpoints, inspect the lookup tables and adjust the breakpoints. The warning message indicates that the lookup table has a mask that covers less than half the rows or columns.

Restart Optimization

You can restart the optimization by canceling the feature fill and adjusting the feature fill settings. If you restart the optimization, CAGE initializes at the final values from the current lookup table. Consider:

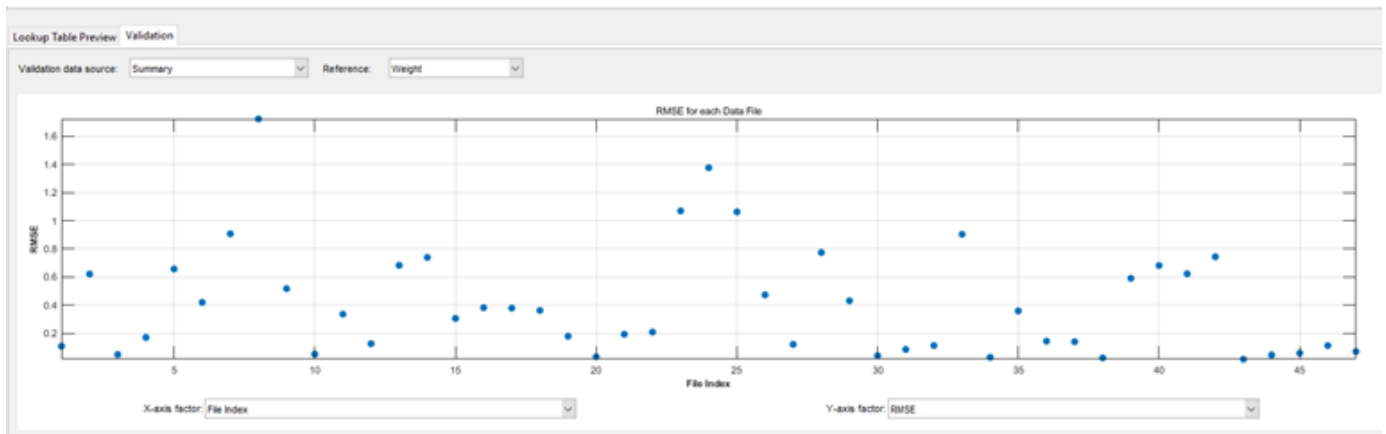
- Switching from Chebyshev to direct table fill when the fill starts to converge to realistic values.
- Changing the step size if the previous run was unstable or if there was not much improvement.
- Adjusting the weight factor to put more emphasis on reducing high errors for specific data files.
- Adjusting the smoothing factor to increase or decrease the table smoothness.

Analyze Results

Use CAGE to analyze the feature fill results.

On the **Feature Filling** pane, the **Validation** tab enables you to:

- View the RMSE for each data file.



- View the full-time trace of the optimization for a data file. Use the **Validation data source** to select a data file. Loading the full-time trace can take a long time for large data files.

On the **Lookup Tables** pane, the extrapolation mask highlights table cells that the data uses. CAGE determines cells that the data uses by examining the derivatives of the feature with respect to table cells for columns containing nonzero entries.

	2592.112	2768.022	2954.327	3148.469	3347.89
0.071	0.472	1.368	2.707	3.954	
-0.082	0.284	1.142	2.476	3.823	
0.022	0.31	1.051	2.242	3.442	
0.509	0.7	1.303	2.311	3.278	
1.006	1.097	1.569	2.404	3.137	
1.251	1.249	1.614	2.314	2.861	
1.181	1.096	1.38	1.9	2.415	
0.887	0.727	0.95	1.6	1.879	
0.497	0.265	0.438	0.69	1.345	
0.078	-0.231	-0.105	0.417	0.846	
-0.394	-0.789	-0.713	0.187	0.337	
-0.986	-1.472	-1.446	-0.904	-0.248	
-1.685	-2.257	-2.274	-1.704	-0.889	
-2.265	-2.906	-2.953	-2.35	-1.386	
-2.223	-2.903	-2.978	-2.385	-1.374	
-0.878	-1.571	-1.718	-1.276	-0.502	

Export Calibration

In CAGE, select **File > Export > Calibration**. You can export the tables in calibration formats, including INCA DCM file and Simple CSV file.

See Also

Related Examples

- “Calibrate a Feature for an Estimator Subsystem in an ECU” on page 5-2

More About

- “Fill Feature Lookup Tables” on page 4-18
- “Import and Export Calibrations” on page 3-39
- “Set Up Models” on page 2-11
- “Set Up Variables and Constants” on page 2-7

Validate Feature Fill

After you fill the feature, use validation data to validate the feature fill.


- 1 In the **Common Tasks** pane, select **Validate**.
- 2 In the Select Data File dialog box, open the file that contains the validation data.
- 3 Use the Data Set Import Wizard to select and assign the data columns.
- 4 Optionally, in the Data Set View dialog box, accept name changes. Click **OK**.
- 5 Optionally, in the Import Validation Data dialog box, click **Yes** to add variables that are not included in the data set to the data.
- 6 Optionally, in the Link Data Set dialog box, link data set factors to variables, models, or tables. Click **OK**.
- 7 In CAGE, in the **Processes** pane, click Feature



- 8 In the CAGE **Validation** tab, review the validation results. Set:

- **Validation data source** to the data source file.
- **Reference** to the model that you want to investigate.
- **X-axis factor** and **Y-axis factor** to the values you want to investigate.

For example, here is validation data for calibrating a feature for an estimator subsystem in an ECU.



See Also


More About

- “Calibrate a Feature for an Estimator Subsystem in an ECU” on page 5-2
- “Data Sets in CAGE”
- “Fill Feature Lookup Tables” on page 4-18

Optimizing Normalizer Breakpoints and Table Values

Optimizing breakpoints alters the position of the table normalizers so that the total square error between the model and the table is reduced.

This routine improves the fit between your strategy and your model. The following illustration shows how the optimization of breakpoint positions can reduce the difference between the model and the table. The breakpoints are moved to reduce the peak error between breakpoints. In CAGE this happens in two dimensions across a table.



Filling Methods

The methods for spacing the breakpoints of your normalizers in CAGE are:

- For one-dimensional tables
 - ReduceError
 - ShareAveCurv
- For two-dimensional tables
 - ShareAveCurv
 - ShareCurvThenAve

ReduceError

Spacing breakpoints using ReduceError uses a greedy algorithm. CAGE:

- 1 Locks two breakpoints at the extremities of the range of values.
- 2 Interpolates the function between these two breakpoints.
- 3 Calculates the maximum error between the model and the interpolated function.
- 4 Places a breakpoint where the error is maximum.

- 5 Repeats steps 2, 3, and 4.
- 6 Locates all the breakpoints, ending the algorithm.

ShareAveCurv and ShareCurvThenAve

Consider calibrating the normalizers for speed, N , and relative air-charge, L , in the preceding MBT model.

In both cases, CAGE approximates the $MBT_{AV}(N, L)$ model, in this case, using a fine mesh.

The breakpoints of each normalizer are calibrated in turn. In this example, these routines calibrate the normalizer in N first.

Spacing breakpoints using ShareAveCurv or ShareCurvThenAve calculates the curvature, K , of the model $MBT_{AV}(N, L)$:


$$K = \sum_{i=1}^{\text{fine mesh}} (MBT_{AV}''(N, L))^{1/2}$$

as an approximation for

$$K = \int_{750}^{6000} |MBT_{AV}''(N, L)|^{1/2} dN$$

Both routines calculate the curvature for a number of slices of the model at various values of L . For example, the figure shown has a number of slices of a model at various values of L .

Model Slices at Various Values of L



- ShareAveCurv averages the curvature over the range of L , then spaces the breakpoints by placing the i^{th} breakpoint according to the following rule.
- ShareCurvThenAve places the i^{th} breakpoint according to the rule, then finds the average position of each breakpoint.

Rule for Placing Breakpoints

If j breakpoints need to be placed, the i^{th} breakpoint, N_i , is placed where the average curvature is^[1]:

$$\int_{750}^{N_i} |MBT_{AV}''(N, L)|^{1/2} dN = \frac{i-1}{j-1} \times K$$


This condition spaces out the breakpoints so that an equal amount of curvature (in an appropriate metric) occurs in each breakpoint interval. The breakpoint placement is optimal in the sense that the maximum error between the lookup table estimate and the model decreases with the optimal convergence rate of $O(N^{-2})$. For reference, consider that equally spaced breakpoints have an order of $O(N^{-1/2})$.

Optimizing Table Values

The Feature Fill Wizard optimizes the table values to minimize the current total square error between the feature values and the model.

This routine optimizes the fit between your strategy and your model. Using **Fill** places values into your table. The optimization process shifts the cell values up and down to minimize the overall error between the interpolation between the model and the strategy.

In this example, the green shaded areas show the error between the mesh model (evaluated at the number of grid points you choose) and the table values.



CAGE evaluates the model over the number of grid points specified in the Feature Fill Wizard, then calculates the total square error between this mesh model and the feature values. CAGE adjusts the table values until this error is minimized, using `lsqnonlin` if there are no gradient constraints. Otherwise, CAGE uses `fmincon` with linear constraints to specify the gradient of the table at each cell.

References

[1] Boor, Carl de. *A Practical Guide to Splines*. New York: Springer-Verlag, 1978, pp 46.

See Also

More About

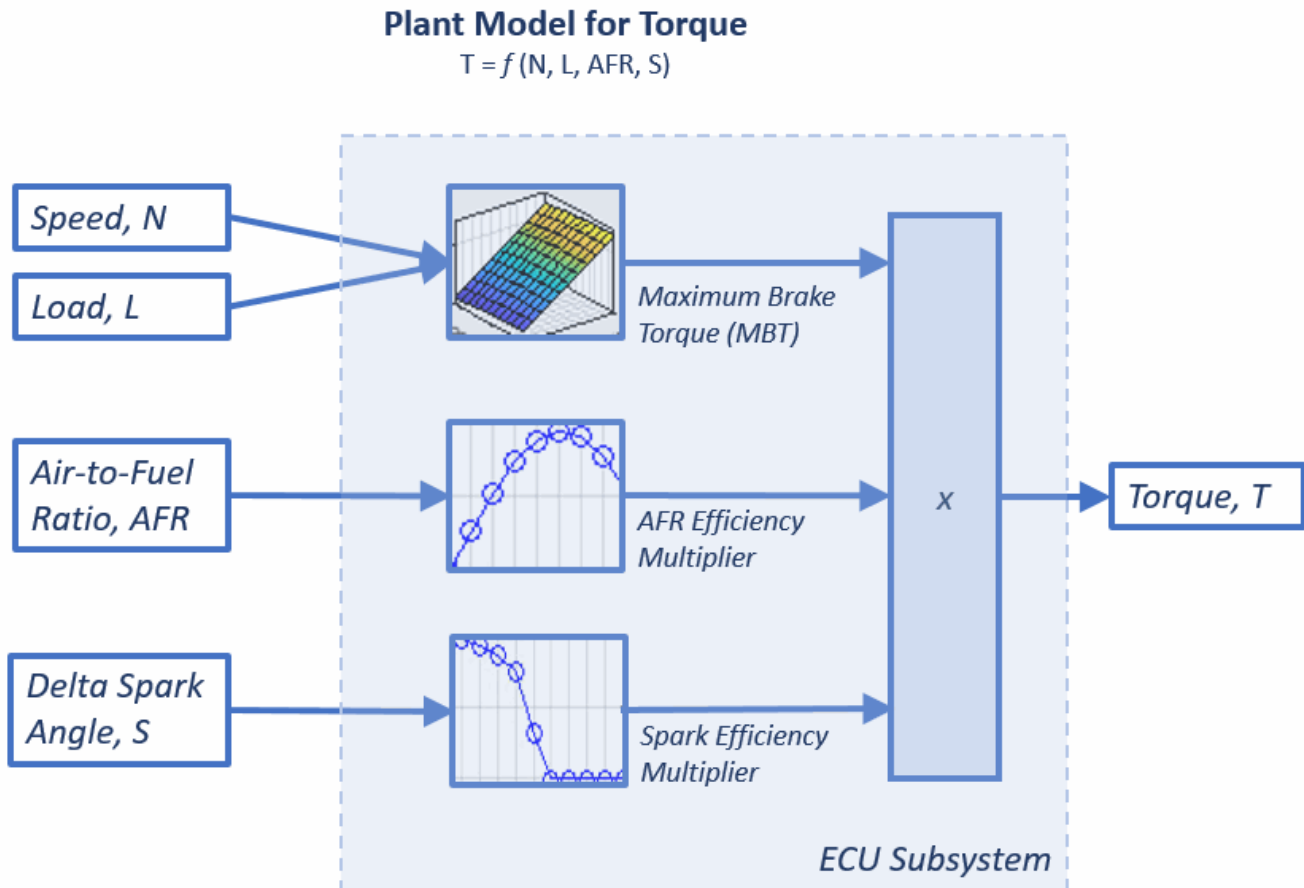
- “About Feature Calibrations and Strategies” on page 4-2
- “Initialize Feature and Calibration Tables” on page 4-12
- “Process for Calibrating and Filling Feature Tables” on page 4-3

Feature Calibration for an ECU

Calibrate a Feature for an Estimator Subsystem in an ECU

This example shows how to calibrate the electronic control unit (ECU) subsystem feature model. The feature model estimates the spark-ignition (SI) engine torque, T , as a function of speed, N , load, L , air-to-fuel ratio, AFR , and delta spark angle, S .

In this example, three lookup tables represent an ECU subsystem.



This table shows a workflow for calibrating an ECU subsystem.

Step	Description
1	<p>“Set Up Feature and Torque Strategy” on page 5-4</p> <p>Create the feature and import the strategy from a Simulink model file. In this example, the strategy represents torque as the product of three lookup tables:</p> <ul style="list-style-type: none"> • Maximum brake torque (MBT) • AFR efficiency multiplier • Spark efficiency multiplier

Step		Description
2	"Import Two-Stage Models" on page 5-6	Use CAGE to import these two-stage statistical models: <ul style="list-style-type: none"> • Torque as a function of engine spark angle, engine load, engine speed, and AFR • Spark at MBT as a function of engine load, engine speed, and AFR
3	"Align Model and Feature Inputs" on page 5-8	To align the statistical model with the feature inputs, create function models.
4	"Fill Feature and Calibration Tables" on page 5-10	Fill the feature so that torque table, T, represents the torque at standard conditions where the engine operates most of the time.
5	"Validate Feature and Export Lookup Tables" on page 5-15	Use validation data to validate the feature fill. Export the calibrated tables.

See Also

More About

- "Set Up Models" on page 2-11
- "About Feature Calibrations and Strategies" on page 4-2

Set Up Feature and Torque Strategy

The feature calibration process calibrates an algebraic collection of lookup tables, or strategy, by comparing the tables to the model.

In this example, you create a feature and import the strategy from a Simulink model file. The strategy represents torque as the product of three lookup tables:

- Speed-load
- Air-to-fuel ratio (AFR)
- Spark angle

1 At the command line, open the CAGE browser.


cage

2 To import the strategy, select **File > Import > Strategy**. Navigate to `matlab\toolbox\mbc\mbctraining`. Open the `torqueFeature` model file.

Model-Based Calibration Toolbox imports lookup tables from the Simulink model.

3 In the MBC Feature Importer dialog box, use the **Lookup Table Preview**, **Feature Details**, and **Model Reference** tabs to review the tables and feature details. Click **OK** to import the feature.


4 In CAGE, on the **Equation** tab, review the `Torque_Feature` feature parsed from Simulink as the output of the algebraic equation of tables.




5 To view the strategy, select **Feature > Edit Strategy**.

Simulink windows open, displaying the `Torque_Feature` strategy and the CAGE libraries.

Close the Simulink windows.



6 To view a simplified feature expression, clear **Feature > Full Strategy Display**.



7 Click the plus next to **Torque_Feature** to expand the Feature tree and observe the tables created by importing the strategy.

- T — Speed and load
- F_A — Air-to-fuel ratio
- F_SPK — Spark angle

The screenshot shows the 'Feature' tree on the left, expanded to show 'Torque_Feature' and its sub-features: 'Norm_L', 'Norm_N', 'F_A', 'Norm_A', 'F_SPK', and 'Norm_SPK'. To the right, a table displays values for variables N and L, and their normalized versions (Norm_L, Norm_N, Norm_A, Norm_SPK) under two different conditions (0.2 and 0.261). The table has a teal background for the data cells.


	N	L	0.2	0.261
Untitled	750		0	0
Torque_Feature	1272.727		0	0
Norm_L	1795.455		0	0
Norm_N	2318.182		0	0
F_A	2840.909		0	0
Norm_A	3363.636		0	0
F_SPK	3886.364		0	0
Norm_SPK	4409.091		0	0
	4931.818		0	0
	5454.545		0	0


Import Two-Stage Models

A Model-Based Calibration Toolbox statistical model is a function of a set of variables. Typically, you construct a model using the Model Browser. In this example, you use CAGE to import these two-stage statistical models:

- Torque as a function of engine spark angle, engine load, engine speed, and AFR
- Spark at maximum brake torque (MBT) as a function of engine load, engine speed, and AFR

The following example uses a model of the behavior of torque with varying spark angles, AFRs, engine speeds, and loads.

- 1 Select **File > Import From Project**. In the CAGE Import Tool, select  to import from a project file.

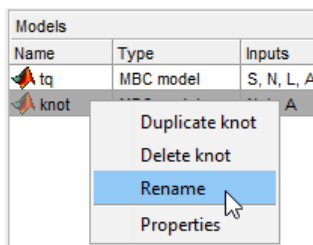


- 2 Browse to `matlab\toolbox\mbc\mbctraining` and open `holliday.mat`. In the CAGE Import Tool, select `tq` and `knot`. Click **Import Selected Items**.


Name	Type	Location
tq	Response	holliday/Two-Stage/tq
tq	Point-by-Point	holliday/Two-Stage/tq/PS22
knot	Datum	holliday/Two-Stage/tq/PS22/knot
knot	Response Feature	holliday/Two-Stage/tq/PS22/knot
max	Response Feature	holliday/Two-Stage/tq/PS22/max
Bhigh_2	Response Feature	holliday/Two-Stage/tq/PS22/Bhigh_2
Blow_2	Response Feature	holliday/Two-Stage/tq/PS22/Blow_2

Close the CAGE Import Tool.

- 3 In the CAGE browser, on the **Models** pane, right-click `knot`. Rename `knot` to `MBT`.



CAGE switches to the **Models** view.




Align Model and Feature Inputs


A model in Model-Based Calibration Toolbox is a function of a set of variables.

In this example, you create two function models to align the statistical model with the feature inputs. Specifically, the function models relate the spark angle to the spark angle retarded from MBT.


- `retard_SPARK` — Maximum brake torque, MBT, minus the spark angle, `S`.
- `abs_SPARK` — Maximum brake torque, MBT, minus the delta spark angle, `delta_S`.

To create function model `retard_SPARK`:

- 1 In the CAGE toolbar, select  to create a function model.





- 2 In the Function Model Wizard:
 - a Enter the formula `retard_SPARK=MBT - S`.
 - b Confirm the model and assigned inputs.




- c Click **Finish**.

To create function model `abs_SPARK`:

- 1 In the CAGE toolbar, select  to create a function model.







- 2 In the Function Model Wizard:
 - a Enter the formula `abs_SPARK=MBT - delta_S`.
 - b Confirm the model and assigned inputs.



- c Click **Finish**.

In the CAGE browser, in the **Models** pane, confirm the model inputs and descriptions in the **Models** view.

Models					
Name	Type	Inputs	Lower Output Limit	Upper ...	Description
 tq	MBC model	S, N, L, A	-Inf	Inf	Created by on
 MBT	MBC model	N, L, A	-Inf	Inf	Created by on :
 retard_SPARK	Function model	MBT, S	-Inf	Inf	MBT-S
 abs_SPARK	Function model	MBT, delta_S	-Inf	Inf	MBT-delta_S

Fill Feature and Calibration Tables

Currently the feature torque table, T , air-fuel-ratio efficiency multiplier table, F_A , and the spark angle efficiency multiplier table, F_SPK , are not filled.

To create the fill settings for the three tables, in CAGE, click Feature Filling



process. To specify the fill settings, use the Feature Fill Wizard.

For this example, to fill the feature so that torque table, T , represents the torque at steady-state engine operating conditions, calibrate the tables in this order.

Fill Order	Table	Fill Description
1	Torque, T	Optimize around conditions which deliver the maximum brake torque (MBT): <ul style="list-style-type: none"> • Spark angle, S, equals the MBT • AFR equals optimal combustion at 14.3
2	Spark, F_SPK	Adjust the torque for lower spark angles by optimizing the delta spark angle, δ_S : <ul style="list-style-type: none"> • Maintain the same speed, N, and load, L • AFR equals optimal combustion at 14.3
3	AFR, F_A	Adjust the torque for different AFRs by optimizing the AFR: <ul style="list-style-type: none"> • Maintain the same speed, N, and load, L • Delta spark angle, δ_S, equals 0

Fill Settings for Torque Table, T

- 1 Select **Feature > New Fill Setting**.
- 2 For bounds, accept the Feature Fill Wizard defaults. Click **Next**.

Choose Tables to Fill

Choose the tables you want to fill and set options on how each table should be filled.

	Table Name	Clear Mask	Extrapolate	Table Bounds	Row Gradient Bo...	Column Gradient ...
<input checked="" type="checkbox"/>	T	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[-Inf, Inf]	[-Inf, Inf]	[-Inf, Inf]
<input type="checkbox"/>	F_A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[-Inf, Inf]	[-Inf, Inf]	
<input type="checkbox"/>	F_SPK	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[-Inf, Inf]	[-Inf, Inf]	

- 3 For set up, set **Filled by** to $tq(s, N, L, A)$. Accept the default input data source and smoothing settings. Click **Next**.

Reference

Feature: Torque_Feature(N, L, A, delta_S)


Filled by: Model

Use percentage error

- 4 Connect input S to function model abs_SPARK. Click **Connect**, then **Next**.

Connect Inputs

Connect model and feature inputs to define fill. Eliminate inputs by connecting to tables. Connect variables describing the same quantities.



- 5 To set the input grid values for:
 - N and L, accept the default values.
 - delta_S, select **Initialize from Normalizer**. In the Select Normalizer dialog box, accept the defaults and click **OK**.
 - A, set value to 14.3


Confirm these input grid values. Click **Finish**.

Set Input Grid Values

Set the grid values you want to use for feature fill.

Name	Value
N	996:364.5454:5006
L	0.2:0.0499:0.7
A	14.3
delta_S	0:5:50

- 6 View the lookup table filling progress. After the filling completes, in the **Lookup Table Preview** tab, view the filled torque table, T.



Fill Settings for Spark Table, F_SPK

- 1 Select **Feature > New Fill Setting**.

- For bounds, set table bounds at $[0, 1]$. Set row gradient bounds at $[-\text{Inf}, 0]$. Click **Next**.

	Table Name	Clear Mask	Extrapolate	Table Bounds	Row Gradient ...	Column Gradient ...
<input type="checkbox"/>	T	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	$[-\text{Inf}, \text{Inf}]$	$[-\text{Inf}, \text{Inf}]$	$[-\text{Inf}, \text{Inf}]$
<input type="checkbox"/>	F_A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	$[-\text{Inf}, \text{Inf}]$	$[-\text{Inf}, \text{Inf}]$	
<input checked="" type="checkbox"/>	F_SPK	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	$[0, 1]$	$[-\text{Inf}, 0]$	

- For set up, set **Filled by** to $\text{tq}(s, N, L, A)$. Accept the default input data source and smoothing settings. Click **Next**.

Reference

Feature: Torque_Feature(N, L, A, delta_S)

Filled by: $\text{tq}(S, N, L, A)$ Model

Use percentage error

- Connect input S to function model abs_SPARK. Click **Connect**, then **Next**.

Connect Inputs

Connect model and feature inputs to define fill. Eliminate inputs by connecting to tables. Connect variables describing the same quantities.

Input connections:

Assigned Input	Input	
	N	
	L	
	A	
	delta_S	
	abs_SPARK(N, ...)	S

Available connections:

Item	Type
retard_SPARK(N,L,A,S)	Function mo...
abs_SPARK(N,L,A,delta_S)	Function mo...
tq(S,N,L,A)	MBC model
MBT(N,L,A)	MBC model

Buttons: **Connect**, **Disconnect**

- To set the input grid values for:
 - N, L, and delta_S, select **Initialize from Normalizer**. In the Select Normalizer dialog box, accept the defaults and click **OK**.
 - A, set value to 14.3


Confirm these input grid values. Click **Finish**.

Set Input Grid Values

Set the grid values you want to use for feature fill.

Name	Value
N	996.364.5454.5006
L	0.2:0.0499:0.7
A	14.3
delta_S	0:5:50

- View the lookup table filling progress. After the filling completes, in the **Lookup Table Preview** tab, view the filled torque table, F_SPK.



Fill Settings for AFR Table, F_A

- 1 Select **Feature** > **New Fill Setting**.
- 2 For bounds, accept the Feature Fill Wizard defaults. Click **Next**.

	Table Name	Clear Mask	Extrapolate	Table Bounds	Row Gradient Bo...	Column Gradient ...
<input type="checkbox"/>	T	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[-Inf, Inf]	[-Inf, Inf]	[-Inf, Inf]
<input checked="" type="checkbox"/>	F_A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[-Inf, Inf]	[-Inf, Inf]	
<input type="checkbox"/>	F_SPK	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	[-Inf, Inf]	[-Inf, Inf]	

- 3 For set up, set **Filled by** to `tq(s, N, L, A)`. Accept the default input data source and smoothing settings. Click **Next**.

Reference
 Feature: Torque_Feature(N, L, A, delta_S)
 Filled by: tq(S, N, L, A) Model
 Use percentage error

- 4 Connect input S to function model `abs_SPARK`. Click **Connect**, then **Next**.

Connect Inputs

Connect model and feature inputs to define fill. Eliminate inputs by connecting to tables. Connect variables describing the same quantities.

Input connections:			Available connections:	
Assigned Input	Input		Item	Type
	<input checked="" type="checkbox"/> N		retard_SPARK(N,L,A,S)	Function mo...
	<input checked="" type="checkbox"/> L		abs_SPARK(N,L,A,delta_S)	Function mo...
	<input checked="" type="checkbox"/> A		tq(S,N,L,A)	MBC model
	<input checked="" type="checkbox"/> delta_S		MBT(N,L,A)	MBC model
abs_SPARK(N, ...	<input checked="" type="checkbox"/> S			

Buttons: Connect, Disconnect

- 5 To set the input grid values for:
 - N and L, select **Initialize from Normalizer**. In the Select Normalizer dialog box, accept the defaults and click **OK**.
 - A, accept the default values

- delta_S, set value to 0

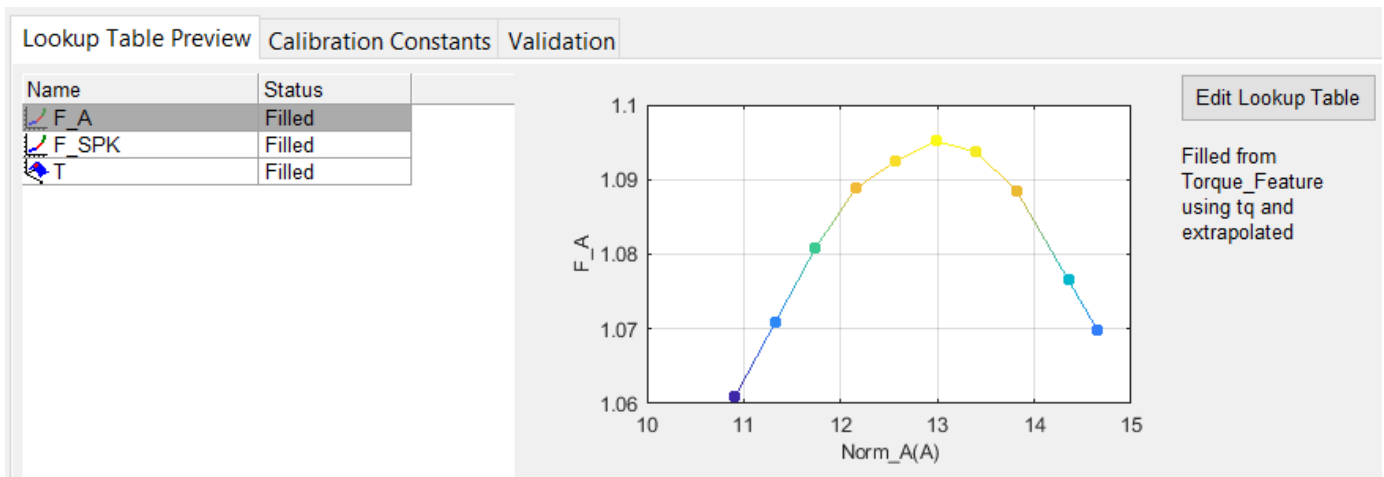
Confirm these input grid values. Click **Finish**.

Set Input Grid Values

Set the grid values you want to use for feature fill.

Name	Value
N	996.364.5454:5006
L	0.2:0.0499:0.7
A	[10.91 11.3256 11.7411 12.1567 12.5722 12.9878 13.4033
delta_S	0

- 6 View the lookup table filling progress. After the filling completes, in the **Lookup Table Preview** tab, view the filled torque table, F_A.



- 7 To view the results of the feature fill, view the **Feature Fill Settings** pane. In this example, the feature filled tables in this order: T, F_SPK, F_A. Review the root mean square error (RMSE) values for each of the table fills.


Feature Fill Settings				
Lookup Tables	Filled by	Grid/Dataset	Smoothing	RMSE
T	tq	Grid(N,L,delta_S)	0	2.755
F_SPK	tq	Grid(N,L,delta_S)	0	1.723
F_A	tq	Grid(N,L,A)	0	1.16

Validate Feature and Export Lookup Tables

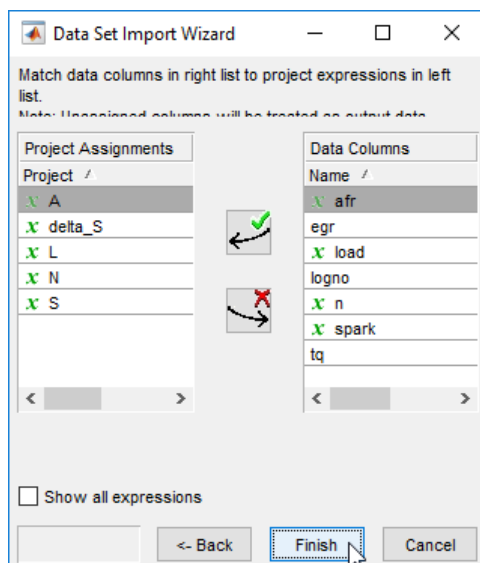
Validate Feature

After you fill the feature, use validation data to validate the feature fill.


- 1 In the **Common Tasks** pane, select **Validate**.
- 2 In the Select Data File dialog box, navigate to matlab\toolbox\mbc\mbctraining. Open torqueValidation.xlsx.
- 3 In the Data Set Import Wizard:
 - a Accept the default data columns to include in the data set. Click **Next**.




- b Accept the default project assignments. Click **Finish**.




- In the Data Set Viewer dialog box, accept the name change. Click **OK**.



- In the Import Validation Data dialog box, click **Yes** to add variables that are not included in the data set to the data.



- In the Link Data Set dialog box, link the `delta_S` data set factor to the `retard_SPARK` function model. Click **OK**.




- In CAGE, in the **Processes** pane, click Feature



- In the **Validation** tab, review the validation results. Set:

- **Validation data source** to `torqueValidation_MBCData`.
- **Reference** to the model that you want to investigate, for example `<tq model>`.
- **X-axis factor** and **Y-axis factor** to the values you want to investigate, for example `tq` and `Error`.

The validation results are similar to these.



Export Calibration Tables

To export your feature calibration tables:

- 1 Select the Torque_Feature.
- 2 Select **File > Export > Calibration > All Items**.
- 3 In the Export Calibration Data dialog box, use **Export to** to specify the exported file type, for example ATI Vision MAT file. Click **OK**.
- 4 Specify a file name and location. Click **Save**.

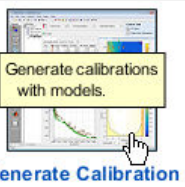
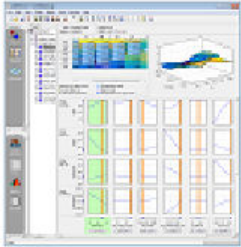
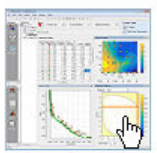
Tradeoff Calibrations

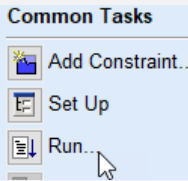
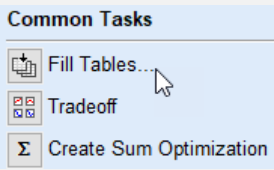
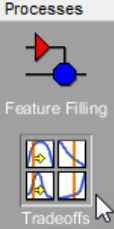
Add New Table

Create a Tradeoff from an Optimization

Typically, there are many different and conflicting calibration objectives. For example, a calibrator might want to maximize torque while restricting nitrogen oxides (NOX) emissions. It is not possible to achieve maximum torque and minimum NOX together, but it is possible to *tradeoff* a slight reduction in torque for a reduction of NOX emissions. A calibrator chooses the values of the input variables that produce the slight torque loss over the values that produce the maximum torque.

Use MBC Optimization app with statistical models created in MBC Model Fitting app to balance competing objectives and fill lookup tables that control engine functions. After you import your models, you can create a tradeoff. The easiest way to automate trading off competing objectives is to create a tradeoff from an optimization. Follow this workflow.


Step		Description
1	"Import Models" on page 6-3	Use the MBC Model Fitting app to import models to CAGE.  Generate calibrations with models. Generate Calibration
2	"Create Lookup Tables and Tradeoff from Model" on page 6-4	Click Lookup Tables and Tradeoff to use the Create Lookup Tables from Model wizard to create the tables and tradeoffs.  Lookup Tables and Tradeoff
3	"Create Optimization" on page 6-7	Click Optimization to use the Create Optimization from Model wizard to create the optimization.  Optimization

Step		Description
4	"Run Optimization" on page 6-9	On the Common Tasks pane, click Run to run the optimization. 
5	"Fill Lookup Tables" on page 6-11	On the Common Tasks pane, click Fill Tables to fill the tables with the optimization results. 
6	"Analyze Results" on page 6-13	On the Processes pane, select Tradeoffs to analyze the tradeoff tables. 


Import Models

Use the **MBC Model Fitting** app to import models to CAGE. For example, import the dual CAM gasoline engine with spark models.


- 1 In the command window, to open the **MBC Model Fitting** app, enter `mbcmodel`.
- 2 In the case studies pane, select Dual CAM gasoline engine with spark optimized during testing.



- 3 On the home page of the MBC Model Fitting app, to open CAGE, select **Generate Calibration**.



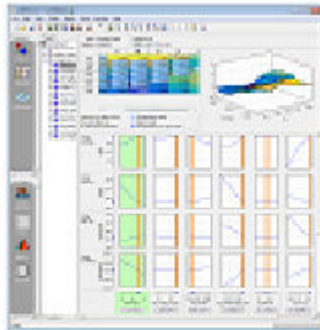
- 4 In the Import Models dialog box, to import the models, select OK. CAGE creates models.



Create Lookup Tables and Tradeoff from Model


In CAGE, use the Create Lookup Tables from Model wizard to create the tables and tradeoff. For example, create tables and a tradeoff for the dual CAM gasoline engine with spark models.

- 1 After you have imported the models for the dual CAM gasoline engine, on the CAGE homepage, select **Lookup Tables and Tradeoff**.




Lookup Tables and Tradeoff


- 2 In the Create Lookup Tables from Model wizard:
 - a Base the new tables on the Torque model. Click Next.



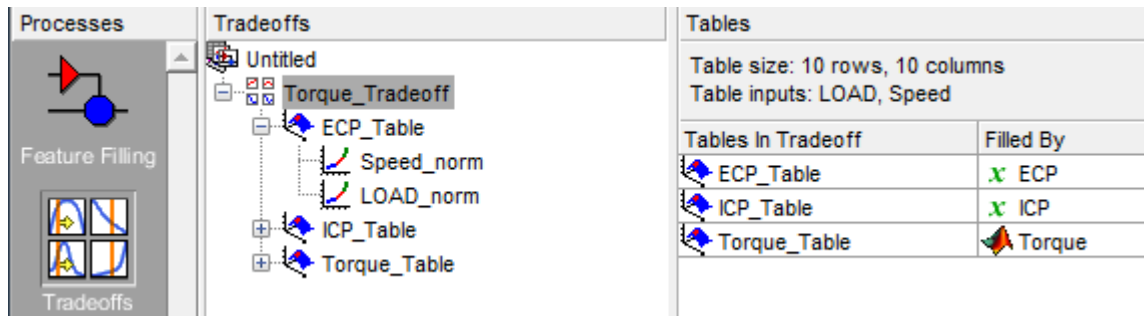
- b Set up the normalizers for the tables. Select Next to accept the defaults.



- c Click Finish to create the tables. By default, the Optimization/Tradeoff fill process is on.



- 3 In the CAGE browser Tradeoffs view, confirm the tables in the tradeoff.




Tables In Tradeoff	Filled By
ECP_Table	✗ ECP
ICP_Table	✗ ICP
Torque_Table	✓ Torque

Create Optimization


In CAGE, use the Create Optimization from Model wizard to create the optimization. For example, create an optimization for the duel CAM gasoline engine with spark.

- 1 After you have created the tables and tradeoffs for the duel CAM gasoline engine, on the CAGE homepage, select **Optimization**.




Optimization

- 2 In the Create Optimization from Model wizard:
 - a Select the Torque model.



- b Set up the optimization. Set these parameters and click Finish.
- **Objective type** to Maximize with Point
 - **Data source** to Breakpoints with Torque_Table (LOAD, Speed)




- c In the CAGE browser Optimization view, confirm the objectives and constraints.

Optimization				
Untitled				
Torque_Optimization				
Objectives				
Name	Description	Type	App	
Torque	Torque(ECP, ICP, LOAD, Speed)	Maximize		
Constraints				
Name	Description	Type	Soft	
Torque_Boundary	Boundary constraint of Torque(ECP, ICP, LOAD, Speed)	Model		

Run Optimization

In CAGE, on the Common Tasks pane, click **Run** to run the optimization. For example, run the optimization for the duel CAM gasoline engine with spark.

- 1 After you have set up the optimization for the duel CAM gasoline engine models, on the Common Tasks pane, click **Run**. View the progress.



- 2 After the optimization completes, view the results.

The screenshot displays the CAGE software interface with several panels:

- Optimization Panel:** Shows 'Current run: 1' and 'Type: Single objective'.
- Results Surface:** A 3D plot showing optimization results with axes for ECP (0-60), Speed (0-4000), and LOAD (0.2-1.2).
- Objective Contours:** A 2D contour plot of Torque vs. ECP, with a color scale from 5 to 15.
- Optimization Results Table:**

Run	Accept	ECP	
1	<input checked="" type="checkbox"/>	24	^
2	<input checked="" type="checkbox"/>	24	
3	<input checked="" type="checkbox"/>	24	
4	<input checked="" type="checkbox"/>	24	
- Constraint Summary:** Shows a constraint named 'Torque_Boun...' with the description 'Boundary constraint of Tor'.
- Common Tasks Panel:** Includes buttons for 'Fill Lookup Tables...', 'Tradeoff', and 'Create Sum Optimization'.
- Summary Information:** States '100 operating points (53 acceptable)' and includes an 'Accept' checkbox.
- Objective Table:**


Objective	Optimization	Table
Torque	12.493	12.493
- Current Operating Point Table:**

LOAD	0.196
Speed	650
- Current Result Evaluation:** Shows 'Optimization Solution' selected.


Fill Lookup Tables

In CAGE, on the Common Tasks pane, click **Fill Lookup Tables** to fill the tables with the optimization results. For example, fill the tables with the duel CAM gasoline engine with spark optimization results.

- 1 After you have run the optimization for the duel CAM gasoline engine models, on the Common Tasks pane, click **Fill Lookup Tables**. In the Lookup Table Filling from Optimization Results wizard, click Next.
- 2 In the Lookup Table Filling from Optimization Results wizard, select the optimization results that you want to fill each table with. Click Next to accept the defaults.




- 3 In the Lookup Table Filling from Optimization Results wizard, select the lookup table filling algorithm. Click Finish to accept the defaults.



Analyze Results

If the tradeoff view is not opened, on the Processes pane, select **Tradeoffs** to analyze the tradeoff tables.



In the tradeoff view, you can:

- Select an optimization for comparison with the current table values.

Option	Goal
Open	Open the optimization results
Run	Set the initial values to the table values and rerun the optimization

Option	Goal
Initialize	Set all tradeoff points to the optimization solution

- Double-click the table lines to set the tables and tradeoff.
 - Orange line updates tables to the tradeoff values.
 - Blue line sets the tradeoff value to the optimization solution at that operating point.
 - Purple line sets the tradeoff value to the table value at that operating point.
- View the tradeoff points on the surface table view. Select the white dots to change the table operating point. Right-click the table surface to edit or print to a figure.

For example, after you have run the optimization for the duel CAM gasoline engine models, on the Processes pane, select **Tradeoffs** to analyze the tradeoff tables.

The screenshot displays the Tradeoffs software interface. On the left, a tree view shows 'Torque_Tradeoff' with sub-items 'ECP_Table', 'ICP_Table', 'Torque_Table', 'Speed_nor', and 'LOAD_nor'. The 'Torque_Table' is selected. The main window shows a table with columns for 'Lookup Table: Torque_Table' and 'Filled by: Torque'. The table contains values for 'LOAD' and 'Speed' at various operating points. A 3D surface plot shows the tradeoff surface. Below the table, there are two 2D line graphs: one for Torque vs. ECP and one for Torque vs. ICP. The Torque vs. ECP graph shows a decreasing trend, and the Torque vs. ICP graph shows a parabolic trend. The current value of Torque is 107. The ECP value is 0.28151367, and the ICP value is 19.6635596.

LOAD	Speed	Torque
0.196	3038.889	15.027
0.323	3009	37.527
0.451	3458	60.339
0.579	3418	83.751
0.707	3954	107.614

Objective	Optimization	Lookup Table
Torque	106.954	106.954

Name	Description	Constraint Value	Left Val
Torque_Boun...	Boundary constr...	7.008e-16	7.008e-

See Also

More About

- “Fill Lookup Tables in a Tradeoff Calibration” on page 6-21
- “Set Up and Perform a Tradeoff Calibration”

Perform a Tradeoff Calibration



A tradeoff calibration is the process of calibrating lookup tables by adjusting the control variables to result in table values that achieve some desired aim.

For example, you might want to set the spark angle and the air/fuel ratio (AFR) to achieve the following objectives:

- Maximize torque
- Restrict CO emissions

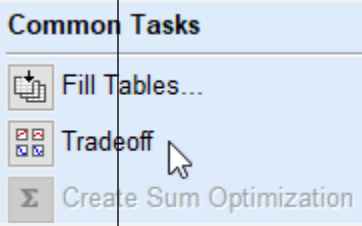
The data in the tradeoff is presented in such a way as to aid the calibrator in making the correct choices. For example, sometimes the model is such that only a slight reduction in torque results in a dramatic reduction in CO emissions.

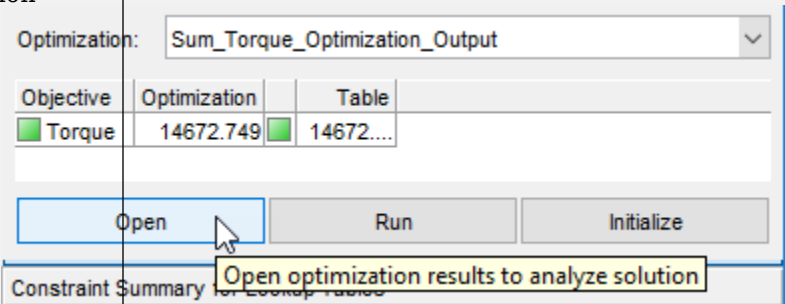
Follow these steps to perform a tradeoff calibration.

Step	
1	“Set Up Variables and Constants” on page 2-7
2	“Set Up Models” on page 2-11
3	“Set Up a Tradeoff Calibration” on page 6-17
4	“Fill Lookup Tables in a Tradeoff Calibration” on page 6-21
5	“Export Calibrations” on page 3-39

You can also use regions to enhance your calibration. See “Control Table Extrapolation Regions” on page 6-31.

In CAGE, you can navigate between the tradeoff plots and the optimization results.

Goal	Option
Navigate from optimization results to tradeoff plots.	In the Common Tasks, select Tradeoff . 

Goal	Option
Navigate from tradeoff plots to optimization results.	<p>In the right-hand-pane, select Open.</p> 

See Also

More About

- “Set Up a Tradeoff Calibration” on page 6-17
- “Fill Lookup Tables in a Tradeoff Calibration” on page 6-21
- “Set Values of Variables” on page 6-25
- “Choose a Table Value at a Specific Operating Point” on page 6-27
- “Control Table Extrapolation Regions” on page 6-31
- “Point-by-Point Model Tradeoffs” on page 6-33
- “Set Up and Perform a Tradeoff Calibration”

Set Up a Tradeoff Calibration

In this section...

“Add a Tradeoff” on page 6-17

“Add Tables to a Tradeoff” on page 6-17

“Display Models in Tradeoff” on page 6-20

A tradeoff calibration is the process of filling lookup tables by balancing different objectives.

Typically, there are many different and conflicting objectives. For example, a calibrator might want to maximize torque while restricting nitrogen oxides (NOX) emissions. It is not possible to achieve maximum torque and minimum NOX together, but it is possible to trade off a slight reduction in torque for a reduction of NOX emissions. A calibrator chooses the values of the input variables that produce this slight loss in torque instead of the values that produce the maximum value of torque.



A tradeoff also refers to the object that contains the models and tables. Thus, a simple tradeoff can involve balancing the torque output while restricting NOX emissions.

After you set up your variable items and models, follow these steps to set up your tradeoff calibration.

Step	
1	“Add a Tradeoff” on page 6-17
2	“Add Tables to a Tradeoff” on page 6-17
3	“Display Models in Tradeoff” on page 6-20

Add a Tradeoff

To add a tradeoff to your session, select **File > New > Tradeoff**. This automatically switches you to the Tradeoff view and adds an empty tradeoff to your session.

An incomplete tradeoff is a tradeoff that does not contain any tables. If a tradeoff is incomplete, it is displayed as  in the tree display. If a tradeoff is complete, it is displayed as  in the tree display.

After you add a tradeoff, you must add tables to your tradeoff.

Add Tables to a Tradeoff

- 1 Add a table by selecting **Tradeoff > Add New Lookup Table** or click  in the toolbar.

You must select the top tradeoff node in the tree display to use the **Tradeoff** menu. This is automatically selected if your tradeoff has no tables yet (it is the only node). You must also add at least three variables (in the variable dictionary) to your project before you can add a table, because CAGE needs a variable to fill the table and two more variables to define each of the two normalizers.

A dialog box opens.

The screenshot shows a 'Table Setup' dialog box with the following fields and controls:


- Name:** Text box containing 'New_2D_Table'
- Rows (Y):** Spinner box set to 10, and a dropdown menu with 'N' selected.
- Columns (X):** Spinner box set to 10, and a dropdown menu with 'L' selected.
- Initial value:** Spinner box set to 0.
- Fill table with:** Text box (empty) and a 'Select...' button.
- Buttons:** 'Clear', 'OK', 'Cancel', and 'Help' buttons.

- 2 Enter the name for the table.


If your tradeoff already contains one or more tables, when you add additional tables they must be the same size and have the same inputs (and therefore have the same normalizers). So if your tradeoff has existing tables, you can only enter the new table name and the initial value.

For the first table in a tradeoff, you must set the normalizer inputs and sizes:

- a Edit the names for the X and Y normalizer inputs (the first two variables in the current variable dictionary are automatically selected here).
 - b Enter sizes for each of the normalizers (Y input = rows, X input = columns)
- 3 Enter an initial value to fill the table cells, or leave this at zero.
 - 4 Click **Select** to choose a filling item for a table. A dialog box opens where you can select from the models and variables in your session.



- a Depending on what kind of input you want, click the radio buttons to display models or variables or both. You can choose to also show items that are filling another table by clearing the check box.
 - b Select the filling item for the table and click **OK**.
- 5 Click **OK** to close the Table Setup dialog box and create the table.

CAGE adds a table node to the tradeoff tree. You can still change the input for the table as follows. Double-click the new table in the list under **Tables In Tradeoff**, or click to select the table (it is selected automatically if it is the only table in the tradeoff) and then click Change Filling Item () in the toolbar.

The Select Filling Item dialog box appears where you can select inputs to fill the table.

- 6 Repeat this procedure for each new table you want to add. Each additional table in the tradeoff must have the same normalizers as the first table, so you do not have to select normalizer inputs and sizes repeatedly. For each new table, you only have to enter the name and initial value.

Add Existing Lookup Tables

- 1 Add a table by selecting **Tradeoff > Add Existing Lookup Tables** or click  in the toolbar.

A dialog box appears where you can select from a list of tables in the current session.



- 2 Select a table and click **OK**. It may be helpful to first select the check box to only show suitable tables that can be added to the tradeoff.

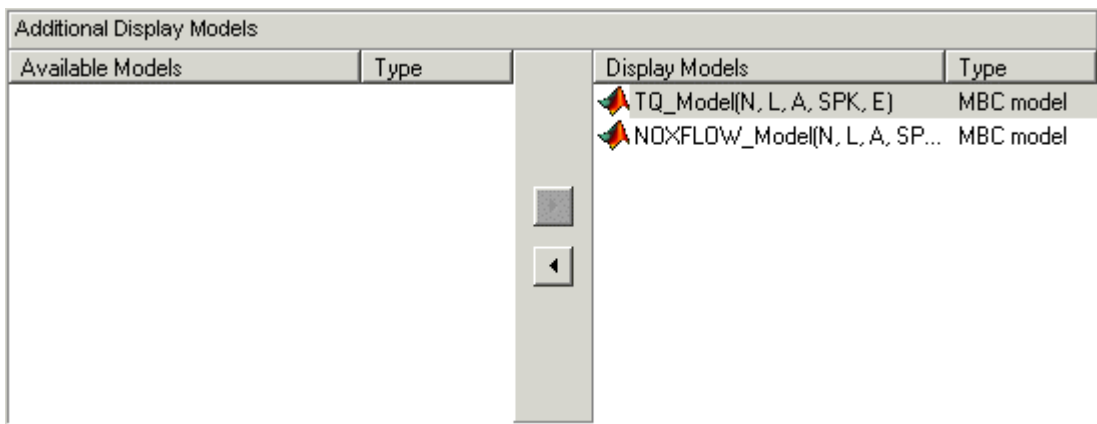
Display Models in Tradeoff

To display models when viewing tables in the tradeoff display:



- 1 Highlight the tradeoff node in the tree.
- 2 From the **Available Models** list, select the one you want to display.

Models that are filling a table are automatically displayed.

- 3 Click  Add Model to Display List in the toolbar or  in the **Additional Display Models** pane to move the selected model into the **Display Models** pane. To quickly add all available models to the display list, click the display button repeatedly.
- 4 Repeat steps 2 and 3 to add all the models you want to the display list.



Remove a Model

- 1 In the **Display Models** list, select the model that you want to remove.
- 2 Click  in the toolbar, or  in the **Display Models** pane, to move the selected model into the **Available Models** pane.
- 3 Repeat until you have cleared all the appropriate models.

Once you have displayed all the models that you want to work with, you are ready to calibrate your tables.

See Also



More About

- “Fill Lookup Tables in a Tradeoff Calibration” on page 6-21

Fill Lookup Tables in a Tradeoff Calibration

Selecting a table node in the tree display enables you to view the models that you have displayed and calibrate that table.

To calibrate the tables, follow these steps.

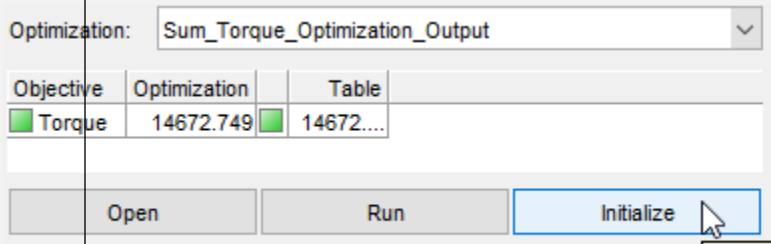
Step	
1	Select the table that you want to calibrate.
2	Highlight one operating point from the table.
3	Set the values for other input variables. See “Set Values of Variables” on page 6-25.
4	Determine the value of the desired operating point. See “Choose a Table Value at a Specific Operating Point” on page 6-27.
5	Click  to apply this value to the lookup table. This automatically adds the point to the extrapolation mask.
6	Repeat the steps to choose values at various operating points.
7	Extrapolate to fill the table by clicking  in the toolbar.
8	Edit table cell values manually by typing values, or right-click to Copy or Paste values.

The graphs colored green indicate how the highlighted table is filled:

- If a row of graphs is highlighted, the table is being filled by the indicated model evaluation (the value shown at the left of the row).
- If the column of graphs is green, the table is being filled by the indicated input variable (shown in the edit box below the column).

You can initialize the entire tradeoff to set the tradeoff points for all operating points. In the CAGE tradeoffs view, after you select a tradeoff table, use these settings.

Goal	Option
Set all tradeoff points to the table values	Inputs > Initialize from Tables

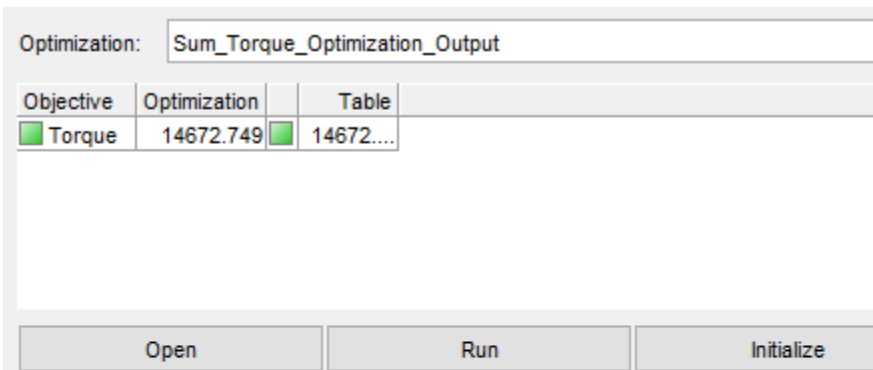
Goal	Option						
Set all tradeoff points to the optimization solution	<p>Either:</p> <ul style="list-style-type: none"> In the toolbar, select Inputs > Initialize from Optimization In the right-hand pane, select Initialize  <p>Optimization: Sum_Torque_Optimization_Output</p> <table border="1"> <thead> <tr> <th>Objective</th> <th>Optimization</th> <th>Table</th> </tr> </thead> <tbody> <tr> <td>Torque</td> <td>14672.749</td> <td>14672....</td> </tr> </tbody> </table> <p>Open Run Initialize</p> <p>Initialize tradeoff and tables from optimization solution</p>	Objective	Optimization	Table	Torque	14672.749	14672....
Objective	Optimization	Table					
Torque	14672.749	14672....					
Clear tradeoff points	Inputs > Clear tradeoff						

In the optimization view, you can select **Fill Lookup Tables**.

You can select an optimization for comparison with the current table values.

Goal	Option
Open the optimization results	Open
Set the initial values to the table values and rerun the optimization	Run
Set all tradeoff points to the optimization solution	Initialize

- For single objective optimizations, the view includes a summary of constraints evaluated at the lookup table values for the current operating point.



Optimization: Sum_Torque_Optimization_Output

Objective	Optimization	Table
Torque	14672.749	14672....

Open Run Initialize

Constraint Summary for Lookup Tables

Name	Description	Constraint Value	Left Value	Right Value
Torque_Boundary	Boundary constr...	-0.178	-0.178	0
KIT1	KIT1(ECP(64), ICP...	-0.354	1.646	2
RF1	RF1(ECP(64), ICP(...	-19.986	5.014	25
TSPEED	TSPEED(ECP(64), ...	-1.526e5	47375.783	2e5
TEXH	TEXH(ECP(64), IC...	-25.631	834.369	860

- For modal optimizations, you can select the mode.

Optimization: BSFC_PointOpt_Output

Objective	Optimization	Table
BSFC	217.008	217.079
PEAKP...	9.42	9.378

Open Run Initialize

Modal Optimization Solutions

Mode	BSFC	PEAKPRE...
<input checked="" type="checkbox"/> 1	217.008	9.42
<input type="checkbox"/> 2	228.315	9.018

- For multiobjective optimizations, the view includes the Pareto view. The blue star indicates the optimization result.

Optimization: BSFC_NOX_Output

Objective	Optimization	Table
BSFC	282.628	280.257
NOXF...	9.878e-3	0.014

Open Run Initialize

The plot displays a Pareto front with green squares and orange triangles. A blue star marks the optimization result at approximately (10, 282.628). The x-axis is labeled NOXFLOW $\times 10^{-3}$ and the y-axis is labeled BSFC.

See Also

More About

- “Set Up a Tradeoff Calibration” on page 6-17

- “Set Values of Variables” on page 6-25
- “Choose a Table Value at a Specific Operating Point” on page 6-27

Set Values of Variables


In this section...

“Set Values for Individual Operating Points” on page 6-25

“Setting Values for All Operating Points” on page 6-26

Set Values for Individual Operating Points

Typically the models that you use to perform a tradeoff calibration have many inputs. When calibrating a table of just one input, set values for the other inputs.



To set values for inputs at individual operating points:

- 1 Highlight the operating point in the lookup table.
- 2 Use the edit boxes or drag the red bars to specify the values of the other variables.

In the preceding example, the spark table is selected (the SPK graph is colored green). Specify the values of AFR (A) and EGR (E) to be used, for example:

- 1 Select the spark table node.
- 2 Click in the edit box for A and set its value to 14.3.
- 3 Click in the edit box for E and set its value to 0.

The default values are the set points of variables, which you can edit in the Variable Dictionary.

Setting Values for All Operating Points

For example, if you are using a tradeoff to calibrate a table for spark angle, you might want to set the initial values for tables of air/fuel ratio (AFR) and exhaust gas recycling (EGR).

To set constant values for all the operating points of one table:

- 1 Highlight the table in the tree display.
- 2 Select one operating point in the table.
- 3 Enter the desired value of the cell.
- 4 Right-click and select **Extrapolation Mask > Add Selection**.

This adds the cell to the extrapolation mask.

- 5 Click  to extrapolate over the entire table.


This fills the table with the value of the one cell.

See Also


More About

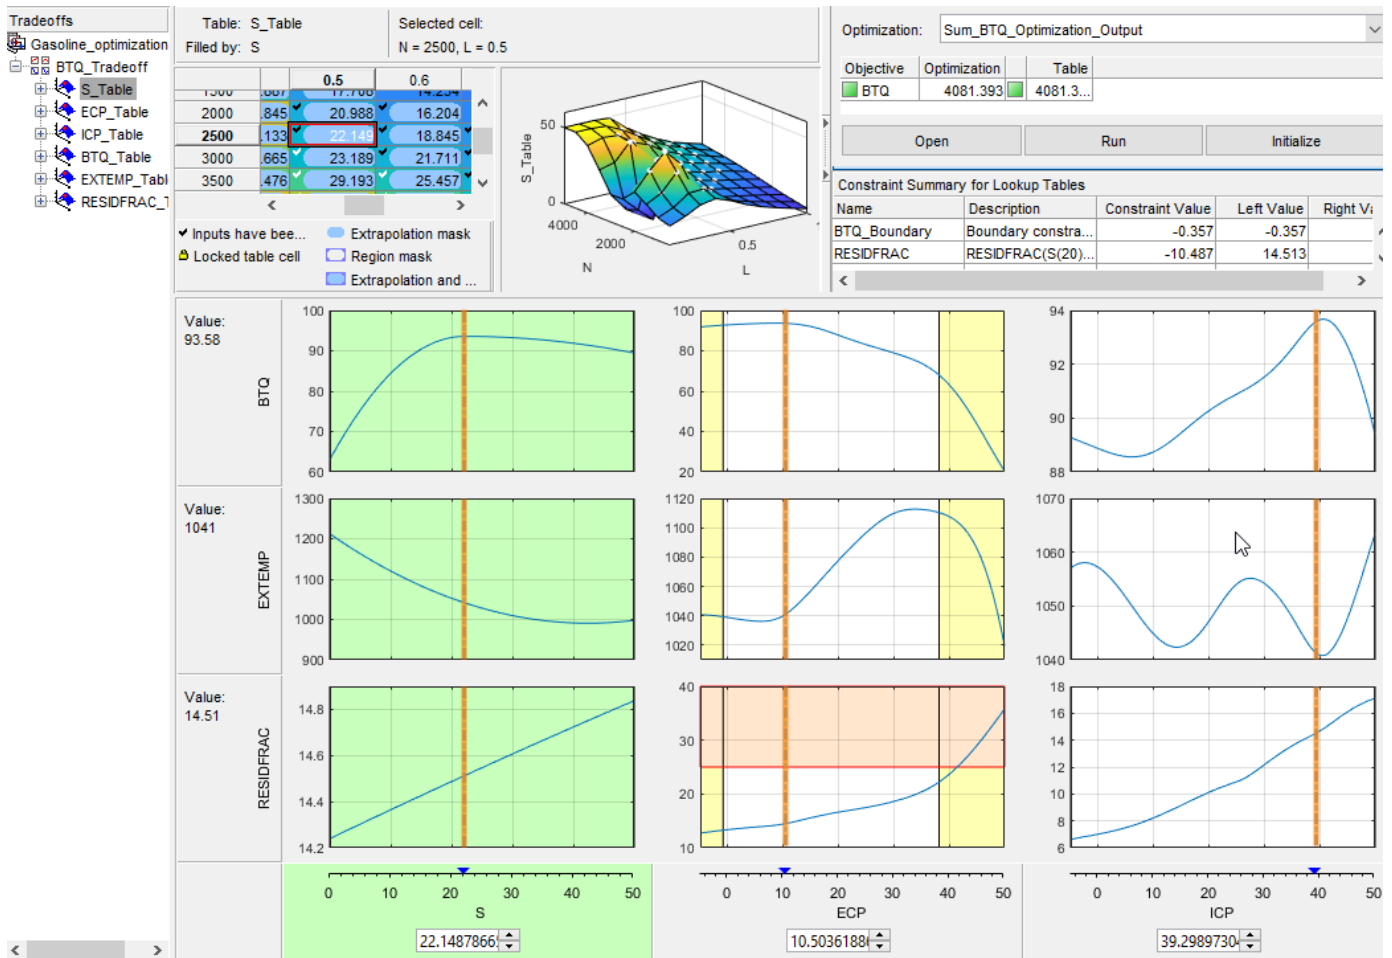
- “Fill Lookup Tables in a Tradeoff Calibration” on page 6-21
- “Choose a Table Value at a Specific Operating Point” on page 6-27

Choose a Table Value at a Specific Operating Point




Performing a tradeoff calibration necessarily involves the comparison of two or more models. For example, in this case, the tradeoff allows a calibrator to check that a value of spark that gives peak torque also gives an acceptable value for the NOX flow model.

- 1 To select a value of an input, do one of the following:
 - Drag the red line.
 - Right-click a graph and select **Find** the minimum, maximum, or turning point of the model as appropriate (also in the toolbar and **Inputs** menu).
 - Click the edit box under the graph and enter the required value.
- 2 Once you are satisfied with the value of your variable at this operating point, apply the value to the table by doing one of the following:
 - Click  (Fill lookup tables with current tradeoff values and save) in the toolbar.
 - Select **Tables > Apply Fill to Table**.



The tradeoff plots show lines for the tradeoff points, table values, and optimization. You can double-click on the lines to set the tables and tradeoff.



- Double-click the orange line to update tables to the tradeoff values.
- Double-click the blue line to set the tradeoff value to the optimization solution at that operating point.

- Double-click the purple line to set the tradeoff value to the table value at that operating point.

Find Maximum, Minimum, or Turning Point of Graphs

Right-clicking a graph enables you to:

- Find minimum of model output with respect to the input variable
- Find maximum of model output with respect to the input variable
- Find turning point of model with respect to the input variable
- Reset graph zooms

Use Zoom Controls on the Graphs

To zoom in on a particular region, shift-click or click with both mouse buttons simultaneously and drag to define the region as a rectangle.

To set the axes limits to the constraint boundary, select **View > Zoom Inside Constraints**. This zooms the tradeoff plots inside the constraint boundaries, so you can explore competing objectives in valid regions only. This option reduces the y-axis scale as the plots do not include regions where the model is extrapolating.

To zoom out to the original graph, double-click the selected graph, or use the right-click **Reset Graph Zooms** option (also in the **View** menu).

Configure Views

Selecting the **View** menu offers you the following options:

Option	Description
Table History	Opens the History display.
Configure Hidden Items	Opens a dialog box that allows you to show or hide models and input variables. Select or clear the check boxes to display or hide items. This is useful if you are trading off a large number of models or models that have many factors.
Display Constraints	Select to toggle constraint displays on and off. Regions outside constraints are shown in yellow on the graphs, as elsewhere in the toolbox.
Display Common Y Limits	Select to toggle a common y-axis on and off for all the graphs. You can also press CTRL+Y as a shortcut to turn common Y limits on and off.
Display Confidence Intervals	Select to display the 99% confidence limits for the models.
Wide Header	Select to toggle graph header size.
Reset Zoom	Use this to reset if you have zoomed in on areas of interest on the graphs. Zoom in by shift-clicking (or clicking both buttons) and dragging. You can also reset the zooms by double-clicking, or by using the right-click context menu on the graphs.

Option	Description
Graph Size	Select from the following options for number and size of graphs: <ul style="list-style-type: none"> • Display All Graphs • Small • Medium • Large

Configure Inputs

Selecting the **Inputs** menu offers you the following options:

Option	Description
Initialize from Tables	Sets all tradeoff points to the table values.
Initialize from Optimization	Set all tradeoff points to the optimization solution.
Clear Tradeoff	Clears tradeoff points
Set to Table Values	Sets the appropriate input value on the graphs to the value in the table.
Reset to Saved Values	Resets all the graph input values to the last saved value.
Apply Fill to Table	Apply the values from the tradeoff graphs to the selected table cell. The corresponding cell in all tables is filled with the appropriate input, not just the cell in the currently displayed table. For example, if you have graphs for spark and EGR inputs, selecting Apply Fill to Table fills the spark table cell with the spark value in the graphs, and the EGR table cell with the EGR value.
Set to Optimization Solution	Set all tradeoff points to the optimization solution.
Select Data Point	Select a data point.
Find Minimum of <i>model vs input factor</i>	Options are only enabled if a graph is selected (click to select, and a blue frame appears around the selected graph). They are also available in the right-click context menu on the graphs. <i>model</i> and <i>input factor</i> are the model and input factor displayed in the currently selected graph, for example, TQ_model versus Spark.
Find Maximum of <i>model vs input factor</i>	
Find Turning Point of <i>model vs input factor</i>	

See Also

More About

- “Set Values of Variables” on page 6-25
- “Control Table Extrapolation Regions” on page 6-31
- “Compare and Revert Lookup Tables” on page 3-17

Control Table Extrapolation Regions

In this section...

“Define a Region” on page 6-31


“Clear a Region” on page 6-31

A region is an area that defines locally where to extrapolate before globally extrapolating over the entire table.

For example, consider filling a large table that has twenty breakpoints for each normalizer by extrapolation. Two problems arise:

- To have meaningful results, set values at many operating points.
- To set values at many operating points takes a long time.

To overcome this problem, you can:


- 1 Define regions within the lookup table.
- 2 In each region, set the values of some operating points.
- 3 Click  to fill the table by extrapolation.

Each region is filled by extrapolation in turn. Then the rest of the table is filled by extrapolation. The advantage of using regions is that you can have more meaningful results by setting values for a smaller number of operating points.

Cells are colored:


- Blue with rounded corners if they form part of the extrapolation mask
- Purple background if they are part of a region
- Both if they are part of the extrapolation mask and part of a region.

Define a Region

- 1 Click and drag to highlight the rectangle of cells in your tradeoff table.
- 2 To define the region, click  in the toolbar, or right-click and select **Regions > Add to Region**, or select the menu option **Tables > Regions > Add to Region**.

The cells in the region are colored with purple backgrounds.

Clear a Region

- 1 Highlight the rectangle of cells in your table.
- 2 To clear the region, click  in the toolbar, or right-click and select **Regions > Remove from Region**, or select the menu option **Tables > Regions > Remove from Region**.

You can clear all regions at once by selecting **Regions > Clear Regions**.

See Also

More About

- “Choose a Table Value at a Specific Operating Point” on page 6-27
- “Point-by-Point Model Tradeoffs” on page 6-33

Point-by-Point Model Tradeoffs

In this section...

“Adding a Point-by-Point Model Tradeoff” on page 6-33

“Calibrating Using a Point-by-Point Model Tradeoff” on page 6-33


There are two types of tradeoff that you can add to your session, a tradeoff of independent models, or a tradeoff of interconnected models: a point-by-point model (or multimodel) tradeoff.

A point-by-point model tradeoff is a specially built collection of models from the Model Browser.

You can build a series of models so that each operating point has a model associated with it. In the Model Browser, you can export models for a point-by-point model tradeoff from the test plan node. The models must be two-stage and must have exactly two global inputs. You can use the point-by-point test plan template to create these models.

Follow this procedure to use a point-by-point model tradeoff to calibrate:

- 1 Import your model and create tables from your point-by-point model.
- 2 Calibrate the tables.
- 3 Export your calibration.

The point-by-point model is only defined for certain cells in the tradeoff tables. These are the operating points that were modeled using the Model Browser part of the toolbox. These cells have model icons in the table. At each of these operating points, you can use the model to trade off, and by doing this you can adjust the value in the table. The point-by-point model is not defined for all other cells in the table and so you cannot use models to tradeoff. You can edit these cells and they can be filled by extrapolation. You trade off values at each of the model operating points in exactly the same way as when using independent models. When you have determined table values at each of the model operating points, you can fill the whole table by extrapolation by clicking .

Adding a Point-by-Point Model Tradeoff


The simplest way to create your point-by-point tradeoff is to:

- 1 Import your model into CAGE.
- 2 Create tables from your point-by-point model.

When you click **OK**, CAGE creates all the tables for the multimodel tradeoff, with breakpoints at the values you have selected.

Calibrating Using a Point-by-Point Model Tradeoff

Each editable operating point in your tables has a model icon in the cell, like this example cell.




These cells have a model defined at that point. You use the display of these models to help you trade off values at these points to fulfill your aims in exactly the same way as when using independent models in "ordinary" tradeoff mode.

- 1 Change input values by dragging the red lines on the graphs or by typing directly into the edit boxes above the graphs. Use the context menu, toolbar or **Inputs** menu to find the maximum, minimum, or turning point of a model if appropriate.
- 2 Look at the model evaluation values (to the left of each row of graphs) and the input variable values (in the edit boxes below the graphs) to see if they meet your requirements.

Remember that the green highlighted graphs indicate how the selected table is filled: if a row is green, the model evaluation value (to the left) fills the table at that operating point; if a column is green, the input variable value (in the edit box below) fills the table.

See the example following; the SPK column of graphs is green, so the value of SPK in the edit box is entered in the table when you click the Apply Lookup Table Filling Values button (📄).



- 3 When you are satisfied with the tradeoff given by the value of your variable at this operating point, you apply this value to the table by pressing **Ctrl+T**, selecting **Tables -> Apply Fill to Table**, or clicking 📄 in the toolbar.
- 4 When you have determined table values at each of the model operating points, you can fill the whole table by extrapolation by clicking 📄.

See Also

More About

- “Set Up a Tradeoff Calibration” on page 6-17
- “Fit a Point-by-Point Model”


Optimization Setup

Using Optimization in CAGE

Overview of Optimization in CAGE

You can use CAGE to solve many automotive optimization problems. For examples of problems you can solve with CAGE, see “Optimization Problems You Can Solve with CAGE” on page 7-3.

To reach the Optimization view, click the **Optimization** button in the left **Processes** pane.



In the Optimization view, you can set up and view optimizations. The view is blank until you create an optimization. When you have optimizations in your project, the left pane shows a tree hierarchy of your optimizations, and the right panes display details of the optimization selected in the tree.

For any optimization, you need one or more models. You can run an optimization at a single point, or you can supply a set of points to optimize. The steps required are

- 1 Import a model or models.
- 2 Set up an optimization.

Optimization functionality in CAGE is described in the following sections:

- The steps for setting up and running optimizations are described in these sections:
 - “Create an Optimization” on page 7-6
 - “Run Optimizations” on page 7-43
- “Optimization Analysis” describes using the optimization output views to analyze your results, fill lookup tables and export results.
- You can define your own optimization functions for use in CAGE. See “Optimization Scripting”.


Parallel Computing in Optimization

By default, the toolbox automatically runs optimizations in parallel if you have Parallel Computing Toolbox. The optimization runs are then executed in parallel. This option can significantly reduce the computation time for larger problems where each run is taking a lot longer than the time it takes to send the problem to another computer.

When you run the optimization, CAGE calls `parpool` to open the parallel pool if necessary, then the optimization runs are executed in parallel. CAGE displays progress messages until the optimization is completed.

Building models in the Model Browser might open `parpool` for you.

If you do not want to run optimizations in parallel, in CAGE, clear **Optimization > Use Parallel**.



Optimization Problems You Can Solve with CAGE

- “Point Optimization Problems” on page 7-3
- “Sum Optimization Problems” on page 7-4

Point Optimization Problems

CAGE provides a flexible optimization environment in which many automotive optimization problems can be solved. These problems can be divided into two main groups, point and sum problems. This section describes point problems.

In a point problem, a single optimization run can determine optimal control parameter values at a single operating point. To optimize control parameters over a set of operating points, an optimization can be run for each point.

Examples of point problems that CAGE can be used to solve are described below:

- Find the optimal spark timing (SPK), intake valve timing (INTCAM) and exhaust valve timing (EXHCAM) at each point of a lookup table whose axes are engine speed (N) and relative load (L).

Optimized values of the control parameters are determined by running the following optimization at each point of the lookup table:

Objective: Maximize engine torque, $TQ = TQ(N, L, SPK, EXHCAM, INTCAM)$

Constraints:

- Residual fraction $\leq 17\%$ at each (N, L) operating point

- Exhaust temperature $\leq 1290^{\circ}\text{C}$ at each (N, L) operating point
- Engine to be operated inside the operating envelope of the engine
- Find the optimal mass of fuel injected (F), rail pressure (P), pilot timing (PT), and main timing (MT) at each point of a lookup table whose axes are engine speed (N) and engine torque (TQ).

Optimized values of the control parameters are determined by running the following optimization at each point of the lookup table:

Objective: Minimize brake-specific fuel consumption, $\text{BSFC} = \text{BSFC}(\text{N}, \text{TQ})$

Constraints:

- Engine out NO_x ≤ 0.001 kg/s at each (N, TQ) operating point
- Engine out Soot emissions ≤ 0.0001 kg/s at each (N, TQ) operating point
- Find the optimum spark timing (SPK) and exhaust gas recirculation (EGR) at each point of a set of operating points defined by engine speed (N), engine load (L) pairs. Optimized values of SPK and EGR are determined by running the following optimization at each point:

Objective: Maximize engine torque, $\text{TQ} = \text{TQ}(\text{N}, \text{L}, \text{SPK}, \text{EGR})$

Constraints: Engine out NO_x ≤ 400 g/hr at each (N, L) operating point

- For a new engine, find out the optimal torque versus NO_x emissions curve for this engine over the operating range of the engine. This is a multi-objective optimization, and CAGE Optimization contains an algorithm (NBI) to solve these problems.

For this example, the optimal torque-NO_x curve is determined by solving the following optimization problem for optimal settings of spark timing (SPK) and exhaust gas recirculation (EGR):

Objectives:

- Maximize engine torque, $\text{TQ} = \text{TQ}(\text{N}, \text{L}, \text{SPK}, \text{EGR})$
- Minimize engine out NO_x = $\text{NO}_x(\text{N}, \text{L}, \text{SPK}, \text{EGR})$

To find out more about solving multiobjective optimization problems in CAGE, see “Set Up Multiobjective Optimizations” on page 7-22.

- For engines with multiple operating modes, find the best operating mode for each operating point. See “Set Up Modal Optimizations” on page 7-25.

To find out more about solving point optimization types of problems in CAGE, see “Create an Optimization” on page 7-6.

Sum Optimization Problems

In a sum optimization, a single optimization run can determine the optimal value of control parameters at several operating points simultaneously. All the control parameters for the operating points are optimized by calling the algorithm once (there is only one call to `fmincon` per run for a sum optimization). This approach contrasts with a point optimization, which has to make a call to the algorithm for every point to find the optimal settings of the control parameters.

- Find the optimal spark timing (SPK), intake valve timing (INTCAM) and exhaust valve timing (EXHCAM) at each point of a look-up table whose axes are engine speed (N) and relative load (L).

Optimized values of the control parameters are determined by running the following optimization once:

Objective: Maximize weighted sum of engine torque, $TQ = TQ(N, L, SPK, EXHCAM, INTCAM)$ over the (N, L) points of a look-up table.

Constraints:

- Difference in INTCAM between adjacent cells is no greater than 5°.
- Difference in EXHCAM between adjacent cells is no greater than 10°.
- At each table cell, residual fraction $\leq 17\%$
- At each table cell, exhaust temperature $\leq 1290^{\circ}\text{C}$
- Find the optimal start of injection (SOI), basefuelmass (BFM), fuel pressure (P), turbo position (TP) and lift of the EGR valve (EGR) at a set of mode points defined by engine speed (N), engine torque (TQ) pairs.

Optimized values of the control parameters are determined by running the following optimization once:

Objective: Maximize weighted sum of brake-specific fuel consumption, $BSFC = BSFC(SOI, BFM, P, TP, EGR, N, TQ)$ over the (N, TQ) mode points.

Constraints:

- Weighted sum of brake-specific NOx must be less than a legislated maximum
- At each mode point, air fuel ratio must be greater than a specified minimum
- At each mode point, turbo speed must not exceed a specified maximum

To find out more about solving these types of problems in CAGE, see "Set Up Sum Optimizations" on page 7-16.

Create an Optimization

In this section...

“Setting Up Optimizations” on page 7-6
 “Creating Optimizations from Models” on page 7-7
 “Tools for Common Optimization Tasks” on page 7-9
 “Optimization Wizard” on page 7-10

Setting Up Optimizations

For any optimization, you need one or more models. Import the model or models you want to optimize into your project. See “Import Models and Calibration Items Using CAGE Import Tool” on page 2-4. After you import your models, you can create an optimization.

Use the following process to set up an optimization:

- 1 Use the wizard for “Creating Optimizations from Models” on page 7-7 to create your optimization.

You can use the wizard to set up any type of optimization: point or sum, single or multiobjective, modal, or multistart optimizations. See “About Point and Sum Optimizations” on page 7-6.
- 2 Add constraints. You can add a boundary model constraint in the wizard. Use the Optimization view to apply other types of constraints (model constraints, linear, ellipsoid, 1-D table, 2-D table, and range). See “Edit Constraint” on page 7-38 for details of all these constraints.
- 3 Choose the points where you want to run the optimization. To do so, you can use the wizard or the Optimization view. In the wizard you can select suitable table breakpoints, data sets, or point-by-point models, or use the variable set points. In the Optimization view you can select the points manually or import them from data sets, tables, or the output of existing optimizations. See “Edit Variable Values” on page 7-30.
- 4 Run the optimization. See “Run Optimizations” on page 7-43.
- 5 View the results. See “Optimization Analysis”.

About Point and Sum Optimizations

You can set up either point or sum optimizations using the Create Optimization From Model wizard.

You can start a common CAGE calibration workflow by creating a point optimization. What is a point optimization? For a point optimization problem, CAGE can determine optimal control parameter values at a single operating point per optimization run. A run is a single call to the optimization algorithm. To optimize control parameters for a set of operating points, CAGE can run an optimization for each point.

You can use a point optimization workflow to find good initial values for a sum optimization. In a sum optimization, a single optimization run can determine the optimal value of control parameters at several operating points simultaneously. For information on the different steps required for setting up sum optimizations, see “Set Up Sum Optimizations” on page 7-16.

Some optimization problems require optimizing more than one objective simultaneously (multiobjective), or multiple solutions per point (modal or multistart). For information on the different steps required for setting up these optimizations, see “Set Up Multiobjective Optimizations” on page

7-22, “Set Up Modal Optimizations” on page 7-25, and “Set Up MultiStart Optimizations” on page 7-28.

For examples of types of optimization problems, see “Optimization Problems You Can Solve with CAGE” on page 7-3

Creating Optimizations from Models

You can use the Create Optimization from Model wizard to set up any type of optimization: point or sum, single or multiobjective, modal, or multistart optimizations. To set up an optimization using a model in your project, from any view in CAGE, use this procedure:

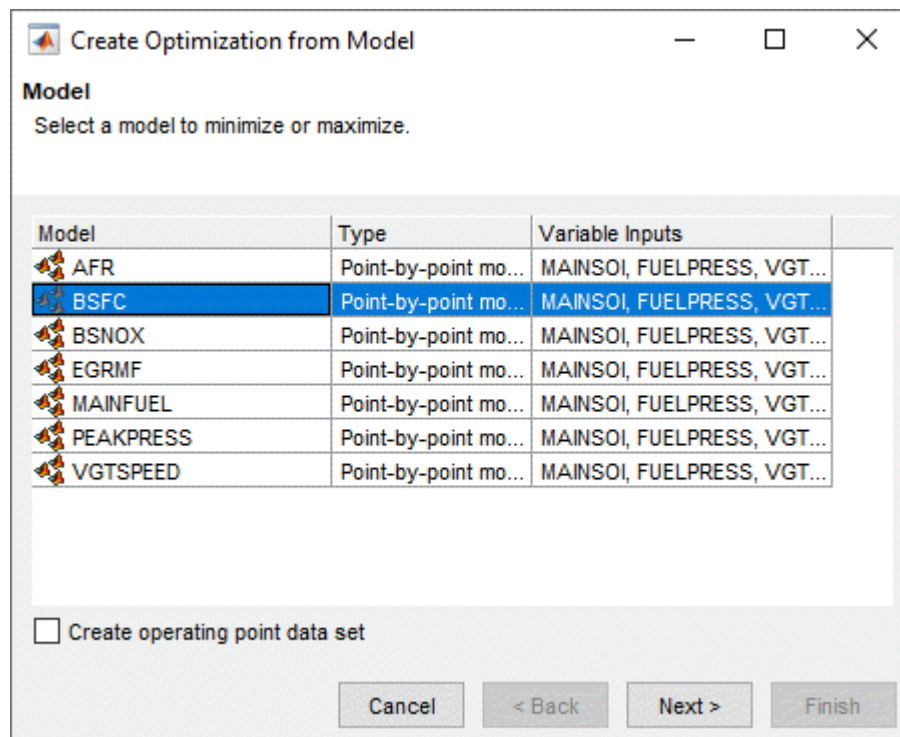
- 1 Select **Tools > Create Optimization From Model** or **File > New > Optimization**.

The **Create Optimization From Model** Wizard appears.

- 2 Select a model to minimize or maximize in the optimization.

If you are viewing a model, then the wizard automatically selects the current model. If you are viewing an optimization or an optimization output node, then the wizard automatically selects the model in the first objective.

If you have point-by-point models as shown in this example, you can optionally select the check box to **Create operating point data set**.



Click **Next**.

- 3 Select the algorithm, objective type, data source for optimization, free variables, and boundary constraint.

Create Optimization from Model

Optimization
Choose optimization type and select free variables to optimize BSFC.

Algorithm: fmincon

Objective type: Minimize Point

Data source: Model operating points

Free variables: 4 selected

Variable	Selected
<input checked="" type="checkbox"/> x MAINSOI	Yes
<input checked="" type="checkbox"/> x FUELPRESS	Yes
<input checked="" type="checkbox"/> x VGTPOS	Yes
<input checked="" type="checkbox"/> x EGRPOS	Yes
<input type="checkbox"/> x SPEED	No
<input type="checkbox"/> x BTQ	No

Add a model boundary constraint

Cancel < Back Next > Finish

- **Algorithm:**
 - Use the default `fmincon` for gradient-based single-objective optimizations.
 - Use `NBI` for multiobjective optimizations. You can set up your additional optimization objectives after you finish the wizard. See “Set Up Multiobjective Optimizations” on page 7-22.
 - Use `gamultiobj` for multiobjective optimizations with more than two objectives and those that are not smooth. Such problems can work better with `gamultiobj` than with `NBI`. You can set up your additional optimization objectives after you finish the wizard. This algorithm is only available if you have the Global Optimization Toolbox product installed. See “Set Up Multiobjective Optimizations” on page 7-22.
 - Use `ga` or `patternsearch` for nongradient-based single-objective optimizations. These algorithms are only available if you have the Global Optimization Toolbox product installed.
 - Use `Modal` optimization with a composite model for selecting the best operating mode for each operating point. See “Set Up Modal Optimizations” on page 7-25.
 - Use `MultiStart` to look for multiple local optimum solutions by running multiple start points for each operating point. See “Set Up MultiStart Optimizations” on page 7-28.
 - If you have a suitable user-defined optimization routine in CAGE, it can appear here. See “Optimization Scripting”.
- **Objective type:** Choose whether you want to Maximize or Minimize your model, and select a Point or Sum objective.

CAGE automatically configures your variable values correctly: a run per point if you select Point, a single run for Sum. See “Set Up Sum Optimizations” on page 7-16.


- **Data source:** Choose the points where you want to run the optimization. You can also set up points in the Optimization view. If you choose to set up points in the wizard, the options depend on the contents of your project and your model type. You can choose to use the variable set points, a data set, table breakpoints, model operating points (the default for point-by-point models), or unique operating points (for a composite model combining point-by-point models).
- **Free variables:** Select the check boxes of the variables you want to optimize from the set of model inputs.
- **Add a model boundary constraint:** Select the check box if you want to constrain the optimization within the boundary model associated with your model.

Click **Finish** to create the optimization.

When you return to the Optimization view you can edit or add constraints and settings and run the optimization.

After you create the optimization, you cannot change the free variables or the algorithm type.

Tools for Common Optimization Tasks



Common tasks are available in the toolbar:

- **Create Tables From Model** — use this wizard to create tables (and optionally a tradeoff) for use with the current optimization. A common workflow is creating tables with the same inputs as your optimization, for filling with optimization results. The wizard automatically selects the model in the first objective (you can also choose any model in your project), and then you can choose which variables and responses to set up tables for. You can then fill these tables with the results from your optimization, and investigate your results in the tradeoff view. See “Create Lookup Tables from a Model” on page 3-2.
- **Add Objective** — Adds an objective to your optimization (if enabled; remember `fmincon` can only have a single objective). Double-click the new objective to open the Edit Objective dialog box, select a model, and set whether to maximize or minimize. See “Edit Objective” on page 7-37.
- **Add Constraint** — Adds a constraint to your optimization. Double-click the new constraint (in the list of constraints) to open the Constraint Editor and set up the constraint. See “Edit Constraint” on page 7-38
- **Import from a data set, import from optimization output, import from table breakpoints, import from table values** — You can use these to populate the Variable Values panes by importing values — See “Edit Variable Values” on page 7-30.
- **Set Up Optimization, Set Up and Run Optimization** — Both these options open the Optimization Parameters dialog box, where you can change optimization settings such as tolerances and number of solutions. When you close the dialog box, the settings are saved (and the optimization runs in the case of Set Up and Run). See “Edit Optimization Parameters” on page 7-44.
- **Run Optimization** — Starts the optimization. See “Run Optimizations” on page 7-43.

Optimization Wizard

For most optimizations, you should use the wizard for “Creating Optimizations from Models” on page 7-7 because it simplifies the setup process.

If your user-defined optimization script defines operating point sets and/or a fixed number of free variables, you must use the Optimization Wizard instead. This is common with Version 2.0 scripts.

For example you may have advanced operating point set requirements, so you need to match the data set variables as part of setting up the free variables; or your optimization problem may involve free variables that are not part of the objective model (e.g., some feasibility problems).

You can use the **Optimization Wizard** to:

- 1 Choose algorithm
- 2 Set up free variables, objectives, and constraints options — “Optimization Wizard Step 2” on page 7-11
- 3 Select free variables — “Optimization Wizard Step 3” on page 7-12

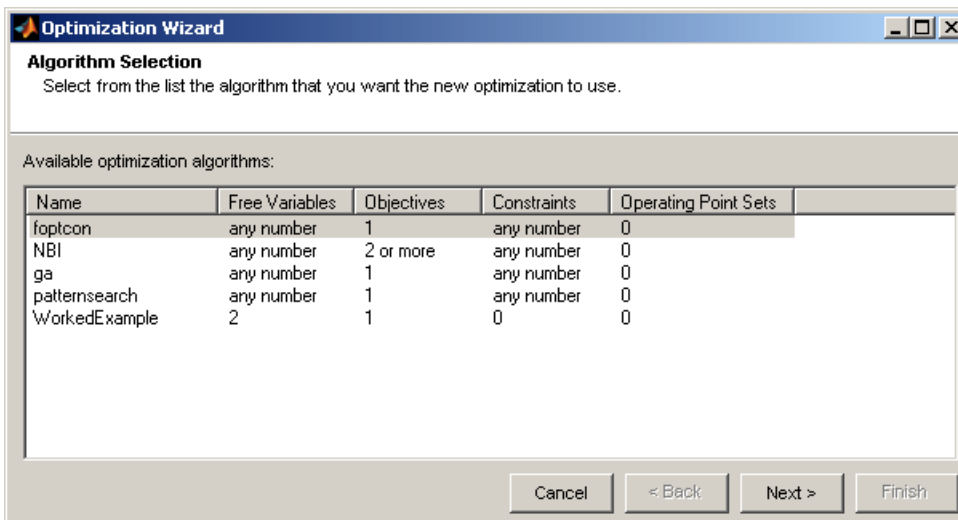
The last 3 steps you can do in the wizard or in the Optimization view:

- 4 Set up objectives — “Optimization Wizard Step 4” on page 7-13
- 5 Set up model constraints — “Optimization Wizard Step 5” on page 7-14
- 6 Set up data sets (user-defined optimizations only) — “Optimization Wizard Step 6” on page 7-14

To use the Optimization Wizard, select **File > New > Custom Optimization**.

This takes you to the Optimization Wizard, which leads you through the steps of choosing the optimization to run, specifying the number of variables to optimize over (unless this is predefined by the function), and linking the variables referenced in the optimization to CAGE variables.

Step 1. First you must choose your algorithm. The first screen of the **Optimization Wizard** is shown below.



The first four algorithm choices in the list are standard routines you can use for constrained single and multiobjective optimization.

- `fmincon` is a single-objective optimization subject to constraints. This function uses the MATLAB `fmincon` algorithm from the Optimization Toolbox™ product.
- `NBI` and `gamultiobj` are multiobjective and can also be subject to constraints.
- `ga`, `gamultiobj`, and `patternsearch` are only available if you have the Global Optimization Toolbox product installed.
 - `ga` stands for Genetic Algorithm, for single-objective optimization subject to constraints. This function uses the MATLAB `ga` algorithm from the Global Optimization Toolbox product. See “Genetic Algorithm” (Global Optimization Toolbox).
 - `patternsearch` is another algorithm for single-objective optimization subject to constraints, from the Global Optimization Toolbox product. See “Direct Search” (Global Optimization Toolbox).
- `Modal optimization` selects the best operating mode for each operating point, and requires a composite model. See “Set Up Modal Optimizations” on page 7-25.

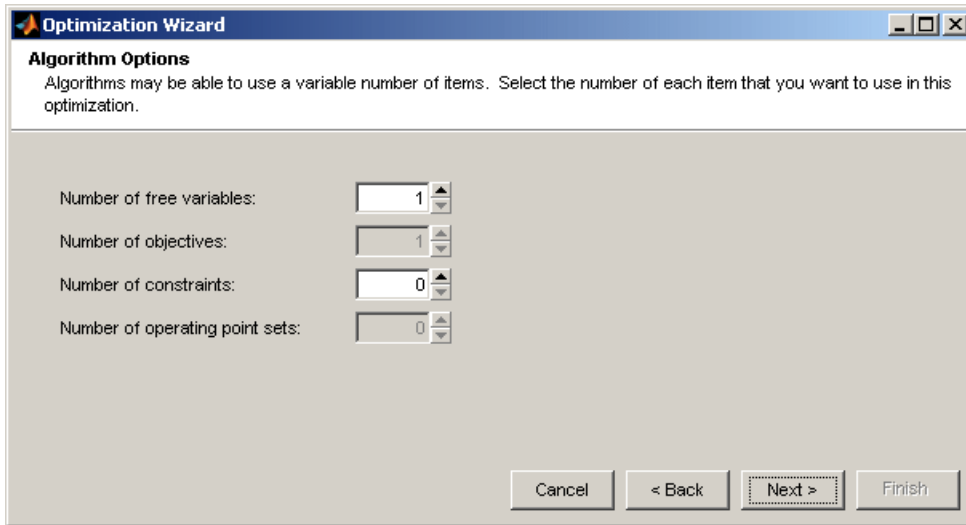
Often these standard routines are sufficient to allow you to solve your optimization problem. Sometimes, however, you might need to write a customized optimization algorithm; to do this you can use the supplied template to modify for your needs. Any optimization functions that you have checked into CAGE appear in this list. See “Optimization Scripting” for information. The Worked Example option is designed to show you how to use the modified template. For step-by-step instructions, see the optimization tutorial section “Create a User-Defined Optimization” on page 9-7.

Note If you choose a user-defined optimization function at step 1, all choices in subsequent steps depend on the settings defined by that function. When writing user-defined optimizations, you can choose to set predetermined algorithm options or allow you to make selections on any subsequent screen of the **Optimization Wizard**.

Optimization Wizard Step 2


Here you select algorithm options for numbers of free variables, objectives, and constraints. The optimization tries to find the best values of the *free* variables. The options available depend on your selected algorithm.

- If in step 1 you select the `fmincon` algorithm and click **Next**, you get the following choices:



The `fmincon` algorithm can only have a single objective, so this control is not enabled. Choose the number of free variables and constraints you require. You can also add constraints later.

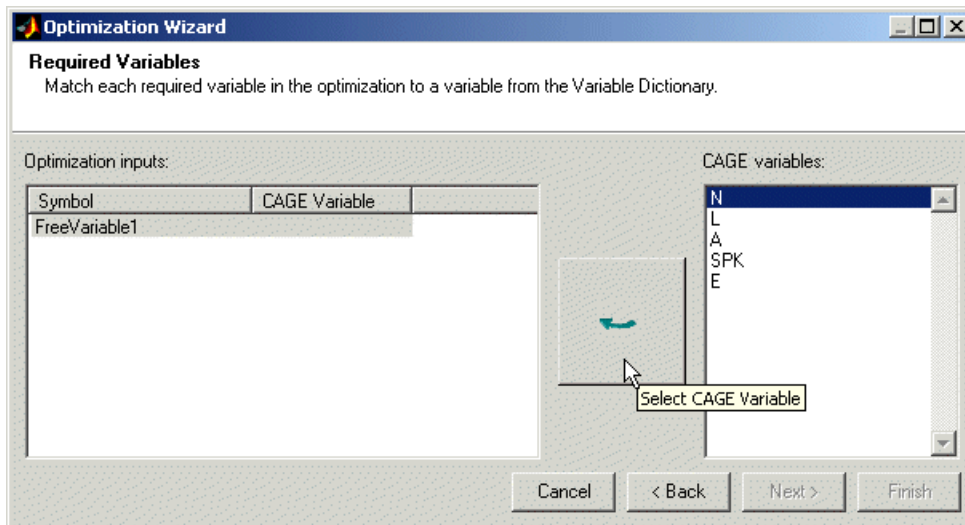
- If in step 1 you select the algorithm NBI, and click **Next**, you see this:



NBI must have a minimum of two objectives, and you can choose as many free variables and constraints as you like. You can add constraints later if necessary.

Optimization Wizard Step 3

You must select variables to link with the free variables used in your optimization.

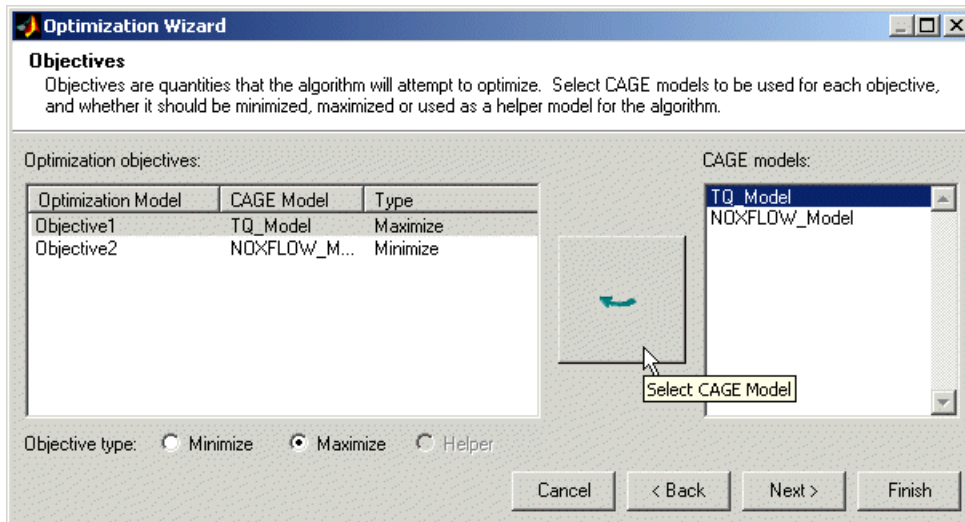


Use this screen to associate the variables from your CAGE session with the free variable(s) you want to use in the optimization. Select the correct pair in the right and left lists by clicking, then click the large button as indicated in the figure.

Once you have assigned your free variables here, you can either click **Next** or **Finish**. This also applies to all later steps in the **Optimization Wizard**.

Optimization Wizard Step 4

You can set up your objectives here or you can set them up at the Optimization view in CAGE. See “Edit Objective” on page 7-37.




Here you can select which models from your session you want to use for the optimization, and whether you want to maximize or minimize the model output. The `fmincon` algorithm is for single objectives, so you can only maximize or minimize one model. The `NBI` or `gamultiobj` algorithms can evaluate multiple objectives. For example, you might want to maximize torque while minimizing NOX emissions. Remember you can also define constraints later, for example, using emissions requirements.

You can also include 'helper' models in your user-defined optimizations, so you can view other useful information to help you make optimization decisions (this is not enabled for NBI or fmincon).

- Click **Next** to proceed to setting up constraints.
- Click **Finish** to complete the **Optimization Wizard** and return to the **Optimization** view. Note you can only set up point objectives in the wizard, but you can also set up sum objectives in the main **Optimization** view. See “Edit Objectives and Constraints” on page 7-37.

Optimization Wizard Step 5

You can use models to define constraint regions that restrict free variables. If you want to use constraints you can select them here, or add them in the Optimization view in CAGE. You can also add other types of constraints in the Optimization view. See “Edit Constraint” on page 7-38.



Select a model for each constraint by selecting a CAGE model and a model constraint and clicking the button to match them up.

For each constraint, either:

- Enter a value in the edit box to define the bound. Select the operator to define whether the optimization output should be constrained to be greater than or less than the value.
- Alternatively, select the option button to use the **Boundary of model** as the constraint.

Optimization Wizard Step 6

If your user-defined optimization allows you to add a data set you can select it on step 6 of the Optimization Wizard. You can use data sets to evaluate models over a different set of operating points during an optimization run. As an example, you could run an optimization at the points (N1, L1), (N2, L2), but an important quantity to monitor and possibly act upon is, say, temperature at points (N3, L3), (N4, L4). You can monitor the temperature at these points by using data sets, to help you select optimization results. You can set up data sets in Step 6 of the wizard or in the Optimization view in CAGE (select **Optimization > Edit Data Sets**).

Data sets are not enabled for fmincon and NBI optimizations.

Click **Finish** to return to the Optimization view in CAGE. Your new optimization appears as a new node in the tree pane on the left, and the setup details appear on the right.

Set Up Sum Optimizations

In this section...

- “Overview of Setting Up Sum Optimizations” on page 7-16
- “Example Problem to Demonstrate Controls for Sum Optimizations” on page 7-17
- “Selecting Scalar Variables” on page 7-18
- “Algorithm Restrictions” on page 7-18
- “Using Application Point Sets” on page 7-20

Overview of Setting Up Sum Optimizations

CAGE can solve sum-type optimizations. These optimizations find the optimal settings of control parameters at several operating points simultaneously. You can use sum optimizations to solve drive-cycle problems where you must apply the constraints across the whole cycle. For example, a constraint such as weighted engine out brake-specific NO_x ≤ 3 g/kWh.

If you have an existing point optimization, you can use a utility to create a sum optimization from your point optimization output. This approach can help you find good initial values for a sum optimization. To create a sum optimization from your point optimization:

- (Recommended) Select **Solution > Create Sum Optimization** . For more details, see “Create Sum Optimization from Point Optimization Output” on page 8-3.
- Select **Optimization > Convert to Sum Optimization** .

If you do not have an existing point optimization, to set up a new sum optimization:

- 1 Use the “Creating Optimizations from Models” on page 7-7 Wizard to create your optimization. You can configure a sum objective in the wizard. CAGE automatically configures your variable values correctly for a sum optimization, defining a single run. See “What Is a Run?” on page 7-17.

You can also configure a sum objective later in the Optimization view. See “Edit Objective” on page 7-37.

- 2 Add constraints:
 - You can add a boundary model constraint in the wizard.
 - To apply other types of constraints you must use the Optimization view. You can apply linear, ellipsoid, 1-D table, 2-D table, and range constraints, and some constraints are specific to sum optimizations—sum constraints and table gradient constraints.

See “Edit Constraint” on page 7-38 for details of all these constraints.

- 3 Choose the points where you want to run the optimization:
 - You can use the wizard to select suitable table breakpoints, data sets, or point-by-point models, or use the variable set points.
 - You can also set up your optimization variable values in the Optimization view. You can enter values manually, or by importing from data sets, tables, or the output of existing optimizations. See “Edit Variable Values” on page 7-30.

For sum optimizations you *must* have a single run, defined by the length controls. CAGE automatically configures your variable values correctly for a sum optimization (defining a single run) when you either:

- Use the **Create Sum Optimization** utility
- Use the **Create Optimization from Model** wizard and then select a sum objective.

If you prefer, you can use the length controls in the Optimization view instead of the wizard. See “Selecting Scalar Variables” on page 7-18.

- 4 (Optional) You can evaluate objectives or constraints over different operating points to those you specified in the optimization. See “Using Application Point Sets” on page 7-20.
- 5 Run the optimization. See “Run Optimizations” on page 7-43.
- 6 View the results (see “Viewing Your Optimization Results” on page 8-14). For descriptions of optimization output specific to sum problems, see “Interpreting Sum Optimization Output” on page 8-54.

What Is a Run?

Sum type optimizations determine optimal settings of operating points simultaneously. Thus, one call to the algorithm determines the optimal settings of the control parameters at each operating point.

In CAGE, a run refers to each call to the optimization algorithm. . You specify the number of runs that you want CAGE to perform with the **Number of runs** control in the **Input Variable Values** pane. For more details, see “Selecting Scalar Variables” on page 7-18.

Example Problem to Demonstrate Controls for Sum Optimizations

The following sections describe the controls and outputs for sum optimizations using the following example problem for illustration.

Say you have created models for torque (TQ), residual fraction (RESIDFRAC) and exhaust temperature (EXTEMP) for a gasoline engine.

The inputs to these models are

- Spark advance, S
- Intake cam timing, INT
- Exhaust cam timing, EXH
- Engine speed, N
- Relative load, L

You need to set up an optimization to calculate optimal settings of S, INT and EXH for the following operating points:

N	L
1000	0.3
1100	0.2
1250	0.31
1500	0.25

N	L
1625	0.18

The objective for this optimization is:

Maximize the weighted sum of TQ over the operating points.

The constraints for this optimization are:

- Constraint 1: EXTEMP \leq 1290°C at each operating point
- Constraint 2: RESIDFRAC \leq 17% at each operating point
- Constraint 3: Change in INT is no more than 5.5° per 500 rpm change in N and 5.5° per 0.1 change in L, evaluated over a 3-by-3 (N, L) table.
- Constraint 4: Change in EXH is no more than 5.5° per 500 rpm change in N and 5.5° per 0.1 change in L, evaluated over a 3-by-3 (N, L) table.

You can use the `fmincon` algorithm in CAGE to solve this problem.

This example is used to explain the controls and outputs in the following sections, “Selecting Scalar Variables” on page 7-18 and “Interpreting Sum Optimization Output” on page 8-54.

See “Algorithm Restrictions” on page 7-18 for details on the optimization algorithm restrictions in CAGE.

Selecting Scalar Variables

At the optimization node, the **Optimization Point Set** pane lists the free and fixed variables. For sum optimizations:

- The **Number of operating points** is the number of points that the optimizer sums during the optimization.
- Use **Select Scalar Variables** to specify the scalar values that do not depend on the operating points.

Algorithm Restrictions

Each run of a CAGE Optimization makes a call to the algorithm you have chosen to use. This algorithm needs to evaluate the objectives and constraints (probably several times) to allow it to determine the optimal settings of the free variables. Optimization algorithms typically have restrictions on the number of objective and constraint outputs they can handle. The following table details the restrictions on the two algorithms provided in CAGE.

Algorithm Name	Objectives	Constraints
fmincon	One output	Any number of outputs
NBI	Two or more outputs	Any number of outputs

When each objective and constraint is evaluated during a run, the number of outputs it returns depends on the maximum number of values of all of its inputs. The following table details the number of outputs each objective type returns as a function of the maximum number of values of all of its inputs.

Objective Type	Maximum Number of Values of All Inputs to the Objective	Number of Outputs	Reason
Point	N	N	A point objective is evaluated at each operating point within a run, and all the values are returned.
Sum	N	One	A sum objective evaluates a model at every operating point and returns one value, which is the weighted sum of the model evaluations.

Similarly, the following table details the number of outputs each constraint type returns as a function of the maximum number of values of all of its inputs.

Constraint Type	Maximum Number of Values of All Inputs to the Constraint	Number of Outputs	Reason
Linear	N	N	These constraints are evaluated at every operating point within a run, and all values are returned.
Ellipsoid	N	N	
1D Table	N	N	
2D Table	N	N	
Model	N	N	
Range	N	0, N or 2N	A range constraint evaluates an expression at each operating point within a run. The constraint returns two values for each point, the distance from the lower and upper bound. In this case 2N outputs are returned. If one of the bounds is infinite, then only the distance to the finite bound is returned for each point, and N outputs are returned. If both bounds are infinite then 0 outputs will be returned.
Sum	N	1	A sum constraint evaluates a model at every operating point and returns the difference between the weighted sum of the model and a bound.
Table	N	≥ 8 (dependent on settings)	A table gradient constraint constrains the gradient of a free variable over a grid. The number of outputs returned depends on the dimensions of the grid.

You can use these three tables to check whether the problem set up satisfies the algorithm restrictions. As an example, the following table checks whether the example problem (detailed in “Example Problem to Demonstrate Controls for Sum Optimizations” on page 7-17) satisfies the restriction of the algorithm chosen to solve it, `fmincon`.

Objective	Maximum Number of Values of All Inputs	Number of Outputs
Weighted sum of TQ over the drive cycle points	5	1 (using the Objective table)
Constraint	Maximum Number of Values of All Inputs	Number of Outputs
EXTEMP \leq 1290°C at each drive cycle point	5	5 (using the Constraint table)
RESIDFRAC \leq 17% at each drive cycle point	5	5 (using the Constraint table)
Change in INT is no more than 5.5° per 500 rpm and 5.5° per 0.1 change in L	5	24 (this value is the number of table gradient constraint outputs generated from a 3-by-3 table)
Change in EXH is no more than 5.5° per 500 rpm and 5.5° per 0.1 change in L	5	24 (this value is the number of table gradient constraint outputs generated from a 3-by-3 table)

Thus, the example problem has 1 objective output and 58 constraint outputs. This satisfies the restrictions of the `fmincon` algorithm and so the algorithm can be used.

Using Application Point Sets

You can use application point sets to evaluate constraints and objectives at different operating points than those specified in the optimization. You can only use application point sets with sum optimizations.

This can be useful for some problems, for example:

- Your calibration problem requires consideration of several drive cycles, defined at different operating point sets. For example, it is common to have different drive cycles for performance and emissions.
- You want to apply some constraints only at a subset of the optimization points (sometimes called “multiregion problems”). For example, there are often different constraints to consider at full load.
- Your full operating point set of interest is large and optimizing at every point is slow. You can run an optimization at a subset of points, and evaluate interpolated results across an application point set.
- You may need to evaluate point-by-point models at different operating points to the points where the models are defined.

To use an application point set for evaluating an objective or constraint:


- 1 Right-click the objective or constraint, and select **Select Application Point Set**.

The Select Operating Point Variables dialog box appears.

- 2 Select a pair of variables to use in application point sets. The variables must be fixed variables in your optimization. You only select variables once per optimization. Click **OK**.

The Select Application Point Set dialog box appears.

- 3 Select an application point set. You can choose a data set or a **New** subset of the optimization points. To select a subset of points, you can use the check boxes or click points on the plot.



- 4 View the plot displaying the application points and optimization points. CAGE extrapolates the optimization results to evaluate the objective or constraint at the application points.
- 5 Click **OK**.

To see example plots that illustrate how CAGE uses application point sets, enter `mbcAppPointSetDemo` to load the example project file `mbcAppPointSetDemo.cag` into CAGE. Run the optimizations in the project to view the plots.

Set Up Multiobjective Optimizations

In this section...

“Overview of Setting Up Multiobjective Optimizations” on page 7-22

“About the gamultiobj Algorithm” on page 7-23

“About the NBI (Normal Boundary Intersection) Algorithm” on page 7-23

Overview of Setting Up Multiobjective Optimizations

CAGE Optimization contains two algorithms (NBI and `gamultiobj`) to solve multiobjective optimization problems. For example, you could use the NBI type of optimization to determine the optimal torque versus NOx emissions curve for an engine over the operating range of the engine. To solve this problem you must define two competing optimization objectives, to maximize torque while minimizing NOx emissions.

To set up a new multiobjective optimization:

- 1 Use the wizard for “Creating Optimizations from Models” on page 7-7 to create your optimization. You can configure one of your objectives in the wizard. You must select either the NBI algorithm or the `gamultiobj` algorithm to solve a multiobjective optimization.

Use `gamultiobj` for multiobjective optimizations with more than two objectives and those that are not smooth. Such problems can work better with `gamultiobj` than with NBI. The `gamultiobj` algorithm is only available if you have the Global Optimization Toolbox product installed.

When you select a multiobjective algorithm, the wizard automatically creates a second blank objective for you. When you finish the wizard and return to the Optimization view, you can configure the second objective (and add a third if desired).

- 2 You can add a boundary model constraint in the wizard. To apply other types of constraints you must use the Optimization view. You can apply linear, ellipsoid, 1-D table, 2-D table, and range constraints, and some constraints are specific to sum optimizations—sum constraints and table gradient constraints.

See “Edit Constraint” on page 7-38 for details of all these constraints.

- 3 You can use the wizard to choose the points where you want to run the optimization. You can select suitable table breakpoints, data sets, point-by-point model operating points, or use the variable set points. You can also set up your optimization variable values in the Optimization view. See “Edit Variable Values” on page 7-30. You can enter values manually or import them from data sets, tables, or the output of existing optimizations.
- 4 Run the optimization using the procedure for “Run Optimizations” on page 7-43.

Click **Run Optimization** in the toolbar to run the optimization with the default settings.

- 5 View the results (see “Viewing Your Optimization Results” on page 8-14). For descriptions of optimization output specific to multiobjective problems, see “Tools for Optimizations with Multiple Solutions” on page 8-38 and “Analyzing Multiobjective Optimization Results” on page 8-51.

About the gamultiobj Algorithm

The `gamultiobj` algorithm uses the `gamultiobj` function from the Global Optimization Toolbox product, and is only available if you have the Global Optimization Toolbox product installed.

For details on the `gamultiobj` function, see “Multiobjective Optimization” (Global Optimization Toolbox) in Global Optimization Toolbox. For CAGE options, see “gamultiobj Optimization Parameters” on page 7-49.


To analyze results, see “Tools for Optimizations with Multiple Solutions” on page 8-38 and “Analyzing Multiobjective Optimization Results” on page 8-51.

About the NBI (Normal Boundary Intersection) Algorithm

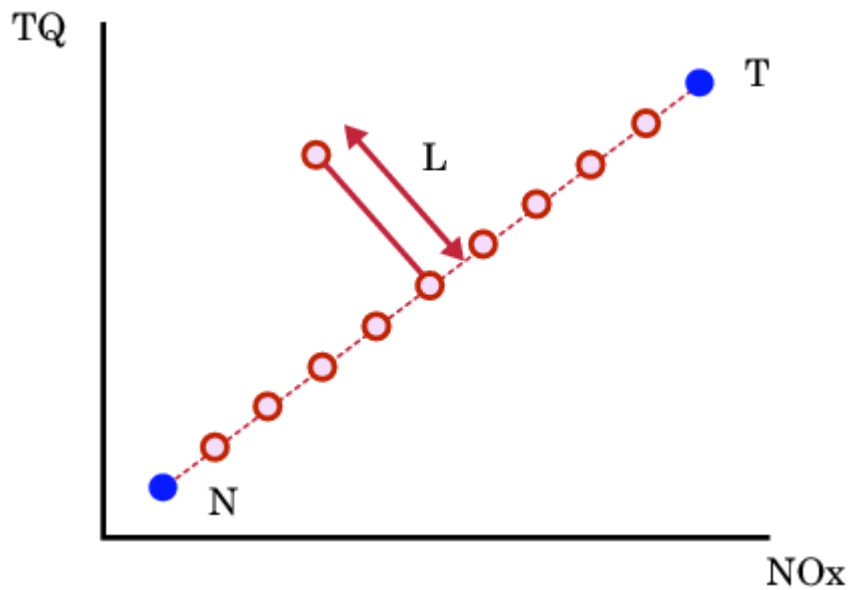
To understand the options for the NBI algorithm, some limited understanding of the algorithm is required. For more information on the NBI algorithm, see the following reference:

Normal-Boundary Intersection: A New Method for Generating the Pareto Surface in Nonlinear Multicriteria Optimization Problems, I. Das and J.E. Dennis, *SIAM J. on Optimization*. 8(3), 631-657 (1998).


The NBI algorithm is performed in two steps. The first step is to find the global of each objective individually. This is called the shadow minima problem, and is a single-objective problem for each objective function. The MATLAB routine `fmincon` is used to find these. Once these are found, they can be plotted against each other. For example, consider an NBI optimization that simultaneously maximizes TQ and minimizes NOx emissions. A plot against each other might resemble the following.



The second step is to find the “best” set of tradeoff solutions between your objectives. To do this, the NBI algorithm spaces N_{pts} start points in the $(n-1)$ hypersurface, S , that connects the shadow. In the above example, S is the straight line that connects the points N and T . For each of the N_{pts} points on S , the algorithm tries to maximize the distance along the normal away from this surface (this distance is labeled L in the following figure). This is called the NBI subproblem. For each of the points, the NBI subproblem is a single-objective problem and the algorithm uses the MATLAB `fmincon` routine to solve it. This is illustrated below for the TQ-NOX example.



The figure above shows spacing of the points along the (n-1) surface. The algorithm tries to maximize the distance L along the normal away from the surface. The following figure shows the final solution found by the NBI algorithm.



To see how the NBI settings are used in the Optimization Parameters dialog box, see "NBI Optimization Parameters" on page 7-45.

Set Up Modal Optimizations

In this section...

“What Is Modal Optimization?” on page 7-25

“Workflow for Modal Optimization” on page 7-25

“Creating Modal Optimizations” on page 7-26

“Adding Extra Objectives to Modal Optimizations” on page 7-26

What Is Modal Optimization?

You can use the modal optimization algorithm to produce optimal calibrations for engines with multiple operating modes. This algorithm helps you choose the best operating mode for each operating point. The algorithm can optimize an objective for each operating mode and select the best solution automatically. You must use a composite model for these optimizations. See “What Are Composite Models?” on page 2-16

The optimization output views help you view and select the best solution for each operating point.

Workflow for Modal Optimization

To find the best mode at each operating point, use the following workflow:

- 1 Create a composite model representing the multiple operating modes. See “Composite Models” on page 2-16.
- 2 Using the Create Optimization From Model wizard, select the Modal optimization algorithm. The wizard then selects the mode input as a free variable.
- 3 (Optional) Add additional objectives in the Optimization view. You can optimize with one objective, and use the other objectives to explore the results with graphical comparisons. You can optionally choose a different objective to use to select the best mode.
- 4 After running the optimization, use the optimization output graphical tools to view the solutions and decide which mode you want to use for each operating point.

The modal optimization algorithm tries to automatically select the best mode for each operating point. Use the optimization output node tools to view all solutions and see which solution is selected. You can change the selections manually if you want. These features are also useful for selecting solutions for multiple objective optimizations, such as NBI.

You can only use modal optimization for point optimization problems. When you have selected the best mode for each operating point, you can create a sum optimization from your modal optimization results. The toolbox automatically creates a sum optimization for you with your selected best mode for each operating point.

To analyze your results, see “Analyzing Modal Optimization Results” on page 8-42.

Examples of Modal Optimizations

Creating Modal Optimizations

To create a modal optimization using a composite model, use the Create Optimization From Model wizard.

- 1 Select **Tools > Create Optimization From Model** (or use the toolbar button).

The **Create Optimization From Model** Wizard appears.

- 2 Select a composite model to minimize or maximize in the optimization. If you are viewing a model, then the wizard automatically selects that current model.

Click **Next**.

- 3 Select the Modal optimization algorithm. The wizard selects the mode input as a free variable. Do not change the Point optimization type.

Specify other optimization information:

- Maximize or minimize
- Data source for optimization
- Free variables
- Boundary constraint

Click **Finish** to create the optimization.

For more information on the options in the wizard, see “Creating Optimizations from Models” on page 7-7.

In the Optimization view, check the Optimization Information pane to confirm your setup. Modal optimizations show the `mbcOSmodalopt` algorithm name and a mode variable, as shown in the following example.

Optimization Information	
Algorithm name	mbcOSmodalopt
Algorithm description	Modal optimization using a single objective optimization subject to constraints
Free variables	CylinderMode, S, ICP, ECP
Operating point variables	None
Mode variable	CylinderMode
Item scaling	off
Distributed runs	off

Your optimization is ready to run, unless you want to change the optimization points, or add constraints or extra objectives. For next steps, see “Adding Extra Objectives to Modal Optimizations” on page 7-26 and “Analyzing Modal Optimization Results” on page 8-42.

Adding Extra Objectives to Modal Optimizations

You can optionally add additional objectives in the Optimization view. You can optimize with one objective, and use the other objectives to explore the results with graphical comparisons. You may want to manually change the selected best mode based on the values of these additional “helper” objectives. You can optionally choose to automatically select the best mode with a different objective.

(Optional) To add extra objectives,

- 1** In the Optimization view, right-click the Objectives pane and select **Add Objective**. Set up the extra objective as described in “Edit Objective” on page 7-37.
- 2** (Optional) You can change which objective to use to select the best mode. By default CAGE uses the first objective to select the best mode, but sometimes you may want to change this. To change this setting,

- a** Select **Optimization > Set Up** or click the toolbar button.

The Optimization Parameters dialog box appears.

- b** Edit the setting **Index to the objective to determine best mode**. The default is 1, so CAGE uses the optimized values of the first objective to select the best mode. Change the index if you want to use a different objective.
- c** Click **OK**.

For next steps, see “Analyzing Modal Optimization Results” on page 8-42.

Set Up MultiStart Optimizations

In this section...
“What Is MultiStart Optimization?” on page 7-28
“Creating a MultiStart Optimization” on page 7-28

What Is MultiStart Optimization?


If you have Global Optimization Toolbox, you can use the `MultiStart` algorithm in CAGE. This algorithm can identify multiple local optimum solutions to help you create smoother tables. Multiple local optimum solutions are often found because of the flat nature of engine responses. The difference in performance between solutions can be small and result in tables that are not smooth enough. The `MultiStart` algorithm tries to identify multiple optimum solutions by running multiple start points for each operating point. You can select the solutions that meet your table smoothness constraints based on the entire table.

The optimization output views help you view and select the best solution for each operating point.

Creating a MultiStart Optimization

- 1 Select **Tools > Create Optimization from Model** (or click the toolbar button) to open the Create Optimization From Model wizard.
- 2 Select a model to minimize or maximize in the optimization. If you are viewing a model, the wizard selects that current model.

Click **Next**.
- 3 Select the `MultiStart` algorithm as shown next. Do not change the `Point` optimization type.



4 Specify other optimization information:

- Maximize or minimize
- Data source for optimization (for example, if you have a suitable table breakpoints to use for optimization points)
- Free variables
- Boundary constraint

5 Click **Finish** to create the optimization.

(Optional) To change the number of start points to run:

- 1** In the Optimization view, select **Optimization > Set Up** or click the Set up optimization toolbar button.

The Optimization Parameters dialog box opens.

- 2** Change the **Number of start points**. This setting determines how many optimizations to run for each point.
- 3** Click **OK**.

Start points is the most common parameter you may want to change. For information on other parameters, see “MultiStart Optimization Parameters” on page 7-48.

Edit Variable Values

In this section...

- “What Are Variable Values?” on page 7-30
- “Define Variables Manually” on page 7-30
- “Import from a Data Set” on page 7-31
- “Import from Output” on page 7-32
- “Import from Lookup Table Breakpoints” on page 7-34
- “Import from Lookup Table Values” on page 7-35

What Are Variable Values?

In the optimization view, the Variable Values panes define the set of operating points for the optimization. If you use the wizard for “Creating Optimizations from Models” on page 7-7, you can choose to set up operating points automatically in the wizard. You can choose to use the variable set points, a data set, table breakpoints, or model operating points (if you have point-by-point models). When you close the wizard, CAGE displays your chosen points in the Variable Values panes. You can use the Variable Values panes to edit your optimization operating points.

You do not have to choose a set of operating points; you can run the optimization at a single point.

Running the optimization requires the selected models to be evaluated (many times over) and hence values are required for all the model input factors. Choose values for the fixed variables in the **Fixed Variables** pane. You chose one or more free variables, so the optimization chooses different values for those free variables in trying to find the best value of the objectives. The initial values for a free variable are shown in the **Free Variables** pane.

To define the set of operating points for the optimization, you can define variables manually, or you can import values from these sources: data set, optimization output, table breakpoints, or table values.

Define Variables Manually

To define values manually:

- 1 In the **Input Variable Values** pane, select the **Number of runs**. New rows appear for both fixed and free variables, all containing the default set point values of each variable. Each row defines an operating point for an optimization run.
- 2 Edit the values in the **Fixed Variables** pane to define the points where you want to run the optimization.
 - You can copy and paste values from other parts of CAGE (existing optimizations or data sets etc.), or from the Help Browser or other documents.
 - You can select **Optimization > Import From Data Set** if you have suitable variables to import.
 - You can select **Optimization > Import From Output** if you have suitable optimization outputs.

An example is shown in the following figure.


Fixed Variables				
Variable:	L	N	A	E
Number of values:	1	1	1	1
1	0.1	1000	12	5
2	0.8	1000	12	5
3	0.1	3000	12	5
4	0.8	3000	12	5
5	0.1	6000	12	5
6	0.8	6000	12	5

- 3 Edit the values in the **Free Variables** pane in a similar way, if you want to define the starting values of the free variables, or you can leave these at the default.
 - For `fmincon` optimizations you can specify various initial starting values per run, see “`fmincon` Optimization Parameters” on page 7-44.
 - If you want to restrict the range of the free variables, you can select **Optimization > Edit Free Variable Ranges**. The default is the range of the variable as defined in the Variable Dictionary.
- 4 Use the right-click context menu to duplicate or delete runs, or select **Fill All Runs** to copy the run values to all other runs.

The Number of Values controls are for sum optimizations. See “Selecting Scalar Variables” on page 7-18.

Import from a Data Set

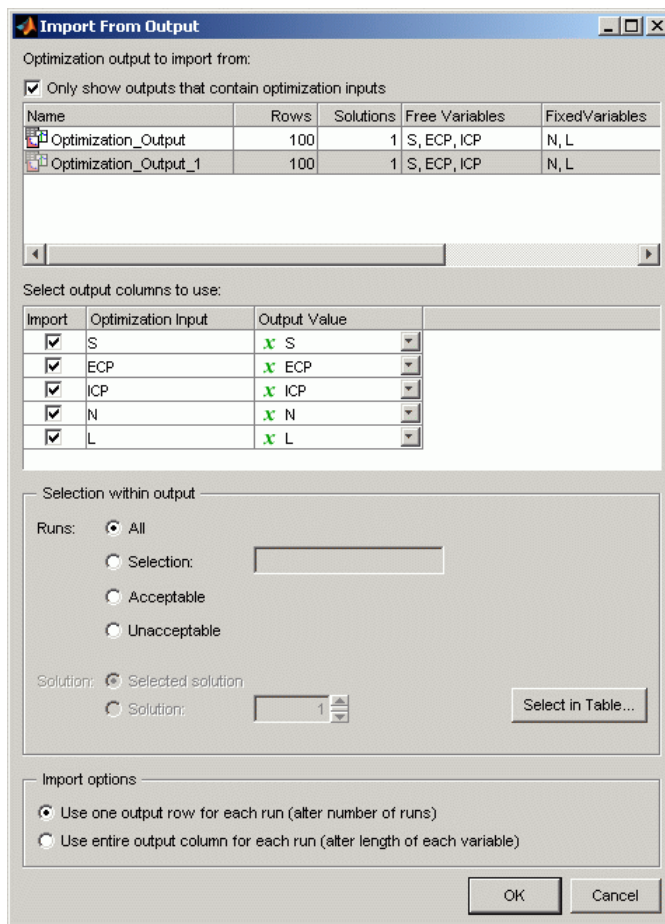
- 1 Select **Optimization > Import From Data Set** (or use the toolbar button) to define the operating points for an optimization from a data set, if you have suitable variables to import. The Import From Data Set dialog box appears.



- 2 Select a data set.
- 3 Select data set columns to import.
- 4 Choose whether you want a run per data set row (**alter number of runs**), or each imported variable to have the same length as the number of data set rows (**alter length**). For information on altering the length of variables (for sum optimizations only), see “Selecting Scalar Variables” on page 7-18.
- 5 Click **OK** to import the variable values.


Import from Output

- 1 Select **Optimization > Import From Output** to import starting values from the output values of a previous optimization. The Import From Output dialog box appears.



- 2 Select the desired optimization output.
- 3 Select the columns from the output you want to import.
- 4 Choose the runs from the optimization output that you want to use. The **Selection within output** controls allow you to choose a subselection. If the number of values per run differs between current inputs and selected outputs, the inputs are altered to match.
 - Select the option button **All** to import all runs.
 - Select the option button **Selection** to import a subset of runs. You can enter a vector specifying the runs you want to import (e.g., 1 3 6:9), or click the button **Select in Table** to open a dialog box and select runs manually.
 - Select the option button **Acceptable** to use only the runs with a selected Accept check box. See “Choosing Acceptable Solutions” on page 8-2. Click the button **Select in Table** to open a dialog box and view or edit the selection.
 - Select the option button **Unacceptable** to use only the runs without a selected Accept check box. Click the button **Select in Table** to open a dialog box and view and edit the selection.
 - For multiobjective optimizations, you can choose to use the selected solutions or a solution number.
- 5 Use the **Import options** buttons to choose whether you want a run per output row (**alter number of runs**), or each imported variable to have the same length as the number of output rows (**alter length**).

If you click the button **Select in Table**, you see the following dialog box.



Highlight cells in the table (**Shift+click**, **Ctrl+click**, or click and drag) to select runs to import.

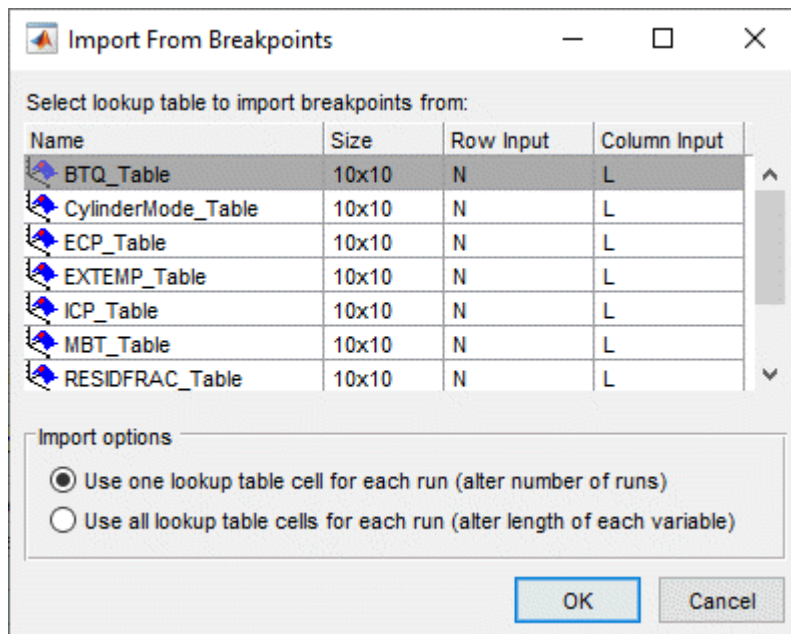
If you chose a subselection on the parent dialog box (e.g., a vector of runs or an acceptable status), the table appears prefiltered with runs selected by those choices. You can filter again for acceptable status on this dialog box: select **Acceptable** or **Unacceptable** from the drop-down list and click the **Select** button.

If there are multiple solutions in the output, you can browse them with the **Solution** controls.

When you are satisfied with the selected runs, click **OK** to return to the Import From Output dialog box. Click **OK** to import the runs.

Import from Lookup Table Breakpoints

- 1 Select **Optimization > Import From Lookup Table Breakpoints** to import starting values from the breakpoint values of a table. The Import From Breakpoints dialog box appears.




- 2 Select the desired table in the list.
- 3 Use the **Import options** buttons to choose whether you want a run per table cell (**alter number of runs**), or each imported variable to have the same length as the number of table cells (**alter length**).
- 4 Click **OK**.

When you click **OK**, values for each table cell are imported into the optimization input variable values pane, e.g., for a 10 by 10 table, 100 starting points are imported.

Import from Lookup Table Values

- 1 Select **Optimization > Import From Table Values** to import starting values from the evaluation of a table. The Import From Table Values dialog box appears.



- 2 For each input you want to import, select the appropriate table from the **Fill Input With** list.

The check box for an input is automatically selected when you select a table for it.

You cannot choose to fill an input with a table that depends on it.

- 3 Click **OK**.

When you click **OK**, your selected optimization inputs are replaced with an evaluation of the table at the current optimization values. Other inputs are not altered.

Edit Objectives and Constraints

In this section...

“Edit Objective” on page 7-37

“Edit Constraint” on page 7-38

“Soft Constraints” on page 7-41

When you create your optimization, you can set up initial objectives and a boundary constraint within the Create Optimization from Model wizard. You can add and edit constraints and objectives in the main CAGE **Optimization** view.

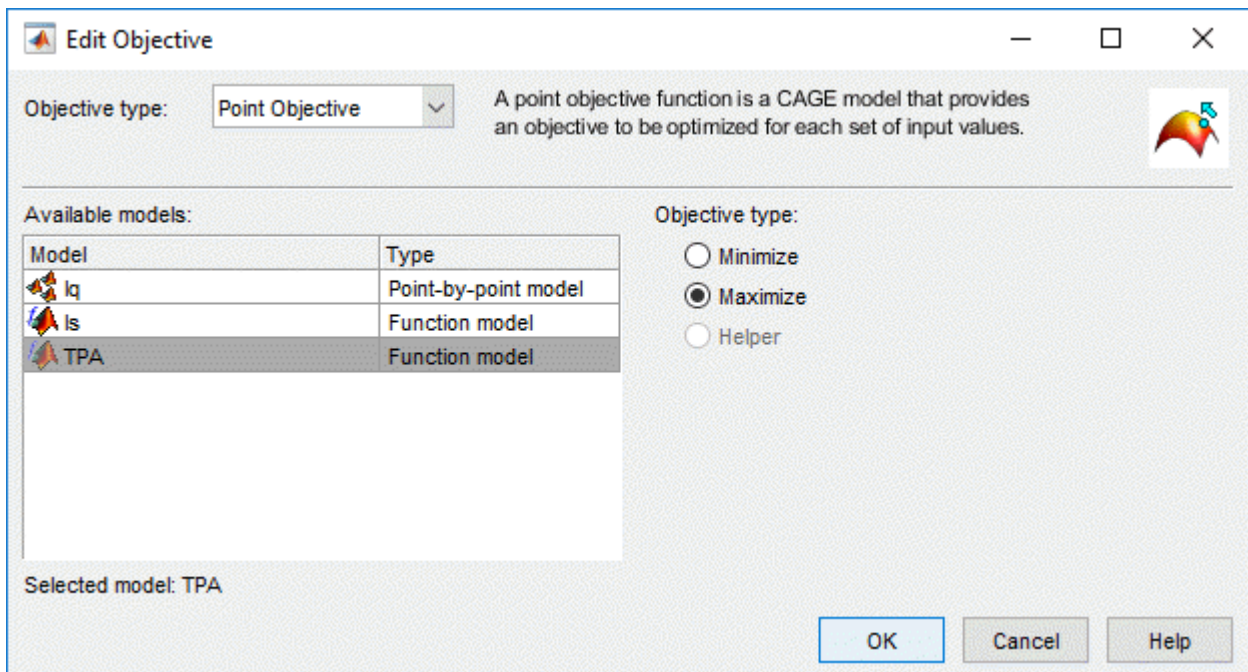
You can run two types of optimizations, point optimizations and sum optimizations. Point optimizations look for the optimal values of each objective function at each point of an operating point set. A sum optimization finds the optimal value of a weighted sum of each objective function. The weighted sum is taken over each point, and the weights can be edited.

Use the Edit Constraint dialog box to set up model sum constraints. You cannot set up sum constraints from the Create Optimization from Model wizard or the Optimization Wizard.

You can also set up linear, 1- and 2-D table, and ellipsoid constraints in the Edit Constraint dialog box.

Edit Objective

Double-click or right-click objectives to open the Edit Objective dialog box.



Objective type Parameter	Description
Point Objective	<p>Point optimizations look for the optimal values of each objective function at each point of an operating point set.</p> <p>The <code>fmincon</code> algorithm is for single objectives, so you can only maximize or minimize one model. The NBI algorithm can evaluate multiple objectives. For example, you might want to maximize torque while minimizing NOX emissions.</p> <p>You can also include 'helper' models in your user-defined optimizations, so you can view other useful information to help you make optimization decisions (this is not enabled for NBI or <code>fmincon</code>).</p>
Sum Objective	<p>Sum optimizations find the optimal value of a weighted sum of each objective function. The weighted sum is taken over each point, and the weights can be edited.</p> <p>Use only for weighted sum optimizations.</p> <p>You can edit weights in the Optimization view, to make certain operating points more important, giving more flexibility to solutions for other points. You can edit the weights in the Fixed Variables pane.</p>

Edit Constraint



You can rename the constraint by editing the **Constraint name** edit box, to aid analysis in the Optimization views. This may be disabled for user-defined optimizations.

Constraint Type Parameter	Description
Model	Optimization constraints.
Range	
Sum Constraint	
Table Gradient	
Linear	These are the same constraints you can apply to designs in the Model Browser part of Model-Based Calibration Toolbox. For more information, see "Define Design Constraints".
Ellipsoid	
1D Table	
2D Table	
	You can select constraint inputs on the additional Inputs tab. Select any variable or model as an input into constraints. The default selects the free variables where possible. Models are treated as nonlinear functions, so if you choose to feed a model into a linear constraint it makes that constraint nonlinear. You are not able to access it as a linear constraint in user-defined optimization scripts.

To improve performance for large optimizations that use these algorithms, consider specifying a **Soft Constraint**.

- fmincon
- modal
- GA
- patternsearch
- gamultiobj
- NBI
- paretosearch
- multistart

During the optimization, Model-Based Calibration Toolbox tries to satisfy the constraint but accepts solutions which are close to the constraint.


Constraints			
Name	Description	Type	Soft Constraint
 TPA_Boundary	Boundary constraint of TPA(ld, Trq, n)	Model	<input type="checkbox"/>
 Is	Is(ld, Trq, n) <= 300	Model	<input checked="" type="checkbox"/>

Model Constraints

To construct a model constraint:

- 1 Select an **Input model** in the left list.
- 2 Use the **Evaluate quantity** drop-down list to choose Evaluation value, Boundary constraint, or PEV value (model prediction error variance) to define your constraint.
- 3 Choose the appropriate option button to either enter a value in the **Constant** edit box, or to select a **CAGE item** from the list of models or variables.
- 4 Select the **Constraint type** operator to define whether the optimization output should be constrained to be greater than, less than, or equal to the constant or item value specified on the right.
- 5 Check the displayed **Constraint description**, and click **OK**.

The model constraint settings are shown in the following figure.



Range Constraints

You can specify an upper and lower bound to constrain expressions (which can be variables, models, or tables). You can specify bounds with constants, vectors, variables, models, or tables.

- 1 Select a CAGE item to constrain on the Bound Expression tab. Use the drop-down menu to switch between variables, models, or tables, and then select the item to constrain. For appropriate models, you can also choose to constrain either the PEV or evaluation value.
- 2 On the Lower Bound tab, select an option button to choose whether to use a constant, vector, or CAGE item to specify the bound.
 - For constants, enter a value.
 - For vectors, you can enter the lower bound for each point in the Input Variable Values pane in the Optimization view after you close the Edit Constraint dialog box.
 - For CAGE items, use the drop-down menu to switch between variables, models, or tables, and then select the item to specify the lower bound. For appropriate models, you can also choose to use either the PEV or evaluation value.
- 3 Specify the upper bound on the Upper Bound tab in the same way as you specified the lower bound on the Lower Bound tab.
- 4 Check the displayed **Constraint description**, and click **OK**.

Sum Constraints

Use these for weighted sum optimizations. Choose a model, constraint bound value, and an operator.

You can have a mixture of point and sum constraints.

Table Gradient Constraints

Table Gradient constraints must be used in sum optimizations. Unless you are using a user-defined optimization, for Table Gradient constraints, use a sum objective. You cannot run point optimizations that contain table gradient constraints.

Table Gradient constraints allow you to constrain the gradient of a free variable or model over a grid of fixed variables.

- 1 Select a free variable or model to constrain.
- 2 Specify one or two fixed variables, and a grid of points either manually or by selecting table axes.
- 3 Enter values in the **Maximum change** and **Axis value change** edit boxes to specify the maximum change in the free variable or model per amount of fixed variable change between cells. For example, enter 5 and 1000 to specify 5-degrees maximum change in cam angle per 1000 rpm.
- 4 To set upper limits or lower limits in a table gradient constraint, specify a two element row vector in the **Maximum change** edit box. e.g., [-5 20].

Use Inf if you only want to specify a lower or upper bound, e.g., enter [0 Inf] to specify a table gradient > 0 , and [-Inf 0] to specify a table gradient < 0 .

- 5 Check the displayed **Constraint description**, and click **OK**.



Soft Constraints

Consider specifying a soft constraint when the:

- Constraint bounds are not well understood.
- Optimization has difficulty finding feasible solutions with hard constraints.

When you specify a soft constraint, Model-Based Calibration Toolbox does not require that the optimization satisfies the constraint. Instead, Model-Based Calibration Toolbox allows for some slack around the constraint. During the optimization, the algorithm satisfies the soft constraint by minimizing the slack and optimizing the objective.

To specify a soft constraint, in CAGE, on the **Constraints** pane, select **Soft Constraint**.

Constraints			
Name	Description	Type	Soft Constraint
 TPA_Boundary	Boundary constraint of TPA(ld, Trq, n)	Model	<input type="checkbox"/>
 Is	Is(ld, Trq, n) <= 300	Model	<input checked="" type="checkbox"/>

For example, follow this general workflow to find feasible optimization solutions using soft constraints.

- 1 Run the optimization with hard constraints.
- 2 Make some constraints soft. For example, for sum optimizations, make the table gradient constraints soft constraints.
- 3 In the Optimization Parameters dialog box, specify a **Soft constraints weighting** to place more or less emphasis on satisfying the constraints. By default, the value is 1. Increasing the value forces the optimization to find feasible solutions while satisfying the hard constraints and

minimizing the slack for soft constraints. If you set the value to inf, the optimization considers all constraints hard constraints. If you set the value to 0, the optimization does not try to satisfy any soft constraints.

- 4 Run the optimization with soft constraints. Review the optimization results. CAGE highlights soft constraints that do not satisfy the hard constraint bound in light blue.

Run	Accept	Torque	Torque ...	KIT1	RF1	TSPEED	TEXH	GradECP	GradICP
(62)			-0.057	-0.735	-18.727	-1.69e5	-36.18	9.447	-8.853
(63)			-0.128	-0.442	-19.392	-1.657e5	-30.951	-18.221	10.671
(64)			-0.359	-0.373	-19.875	-1.528e5	-27.33	-76.553	-1.405
(65)			-0.383	-0.428	-20.193	-1.297e5	-23.44	-105.231	-73.217
(66)			-0.29	-0.457	-20.379	-1.08e5	-19.596	-50.062	-71.696
(67)			-0.214	-0.465	-21.484	-97614.474	-20.753	12.538	-78.315
(68)			-0.113	-0.469	-22.421	-93998.092	-22.179	-19.169	-73.623
(69)			-4.803e-7	-0.465	-22.928	-91285.894	-25.65	-0.954	-42.759
(70)			-0.081	-1.619	-10.92	-1.845e5	-51.281	15.082	-2.488
(71)			-0.24	-1.284	-14.942	-1.778e5	-44.315	15.516	13.215
(72)			-0.339	-0.946	-17.15	-1.724e5	-38.474	-1.458	7.866
(73)			-0.123	-0.573	-18.653	-1.662e5	-32.29	-90.96	-62.687
(74)			-8.396e-3	-0.363	-19.596	-1.616e5	-26.634	-24.477	-2.663
(75)			-0.107	-0.376	-20.151	-1.511e5	-22.25	-19.032	6.406
(76)			-0.335	-0.422	-20.446	-1.313e5	-18.171	-74.084	-10.279
(77)			-0.287	-0.446	-20.478	-1.094e5	-13.261	-92.7	-69.58
(78)			-0.216	-0.473	-21.249	-97003.497	-12.586	-34.457	-73.336
(79)			-0.133	-0.48	-22.074	-89736.521	-11.429	14.599	-73.256
(80)			-0.011	-0.481	-22.826	-85204.851	-12.761	15.217	-69.554
(81)			-0.044	-1.591	-7.656	-1.832e5	-43.759	15.055	-41.653
(82)			-0.193	-1.252	-13.171	-1.764e5	-38.156	11.069	1.242
(83)			-0.343	-0.871	-16.206	-1.707e5	-33.513	14.952	15.525
(84)			-0.331	-0.448	-18.204	-1.636e5	-28.849	-9.571	6.451

- 5 Identify constraints that are not satisfied. Consider making hard constraints soft constraints.

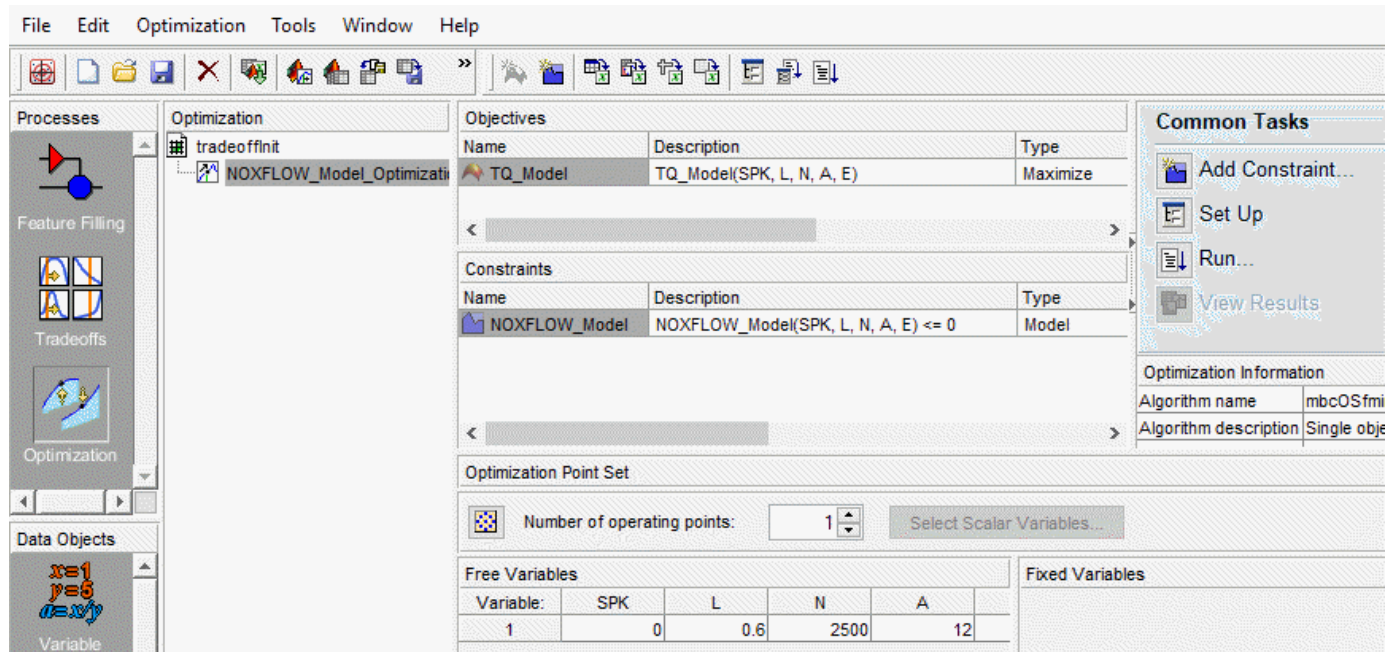
See Also

More About

- “Define Design Constraints”
- “Range Constraint Output” on page 8-23
- “Using Application Point Sets” on page 7-20

Run Optimizations

When you create an optimization, the new optimization appears as a new node in the tree pane on the left. In CAGE, the setup details are on the right.



If your optimization is ready to run, you can click **Run Optimization** in the toolbar to proceed. You may want to edit variable values, constraints, or objectives before running the optimization. If you need to set up any objectives or constraints, **Run** is not enabled. If your optimization is ready to run, you can select **Set Up and Run Optimization** to change algorithm-specific settings such as number of required solutions and tolerances for termination.

- If you click **Set Up**, you can change settings in the Optimization Parameters dialog box.
- If you click **Run** instead, you do not see the optimization settings, but go straight to running the optimization.

You see a progress bar as the optimization proceeds. When the optimization completes, a new Output node appears under your Optimization node in the tree. The view automatically switches to this node so that you can analyze the results.

See Also

More About

- “Edit Objectives and Constraints” on page 7-37
- “Edit Optimization Parameters” on page 7-44
- “Optimization Analysis”

Edit Optimization Parameters

The settings in the Optimization Parameters dialog box are algorithm-specific.

If you edit the settings and want to reset to the defaults, select **Optimization > Reset Parameters**. If you add parameters to user-defined optimization scripts, you might need to use this reset option to make all new parameters appear in the dialog box.

fmincon Optimization Parameters

The fmincon optimization algorithm in CAGE uses the fmincon function in the Optimization Toolbox. fmincon wraps up the fmincon function so that you can use the function for maximizing and minimizing.

Parameter	Description
Constrained optimization algorithm	fmincon function algorithms: Try <code>sqp</code> or <code>interior-point</code> with sum optimizations which are slow or have problems converging. <ul style="list-style-type: none"> <code>active-set</code> (default) <code>sqp</code> — Use for sum optimizations that are slow <code>interior-point</code> — Default option for sum optimizations that have problems converging and optimizations with a large number of constraints
Display	Determines the level of diagnostic information displayed in the workspace. <ul style="list-style-type: none"> (default) <code>none</code> — No information displayed during optimization <code>iter</code> — Displays statistical information every iteration <code>final</code> — Displays statistical information at the end of the optimization
Maximum iterations	Maximum number of iterations allowed.
Maximum function evaluations	Maximum number of function evaluations allowed.
Function tolerance	Termination tolerance on the function value.
Variable tolerance	Termination tolerance on the free variables.
Constraint tolerance	Termination tolerance on the constraint violation.
Barrier update method	Select <code>predictor-corrector</code> (default) or <code>monotone</code> to specify the algorithm to update the barrier parameter for optimizations. <code>predictor-corrector</code> can provide improved performance for sum optimizations.

Parameter	Description
Number of start points	<p>Positive integer, N. Optimization generates (N-1) start points per run in addition to the starting value specified in the Input Variable Values pane.</p> <p>The optimization runs from each of the N start points and chooses the best solution.</p> <p>The N-1 extra start points are generated as follows:</p> <ol style="list-style-type: none"> 1 Generate a 10000 point Halton set design, D, over the free variables. 2 Evaluate the objectives and constraints over D. 3 Return the N-1 feasible points with the lowest objective value. <p>If there are not N-1 feasible points, fill the remaining starting values with the points with the lowest maximum constraint violation.</p> <p>If the specified initial value does not satisfy the constraints, CAGE uses the 10000 search points.</p> <p>Note For point optimization problems, set Number of start points to 1 or 2.</p>
Run from feasible start points only	<p>Select this option to terminate all runs that start with an initial value that does not satisfy the constraints. In the Solution Information pane of the Optimization Output view, the Output message reports the termination.</p>
Soft constraints weighting	<p>Weighting factor that the optimization uses for soft constraints. By default, the value is 1. Increasing the value forces the optimization to find feasible solutions while satisfying the hard constraints and minimizing the slack for soft constraints. If you set the value to inf, the optimization considers all constraints hard constraints. If you set the value to 0, the optimization does not try to satisfy any soft constraints.</p>
Make table gradient constraints feasible	<p>Select this option to make constraints feasible for sum optimizations with table gradient constraints on the free variables.</p> <p>Consider clearing this option if the table gradient constraints are soft.</p>

NBI Optimization Parameters

Use the NBI algorithm for multiobjective optimizations.

Parameter	Description
Tradeoff points per objective pair, N_p	<p>Specify how many tradeoff solutions you want the optimization to find per run.</p> <p>To determine the number of tradeoff solutions between the objectives, N_{pts}, the optimization uses this formula:</p> $N_{pts} = \binom{n + N_p - 2}{N_p - 1}$ <ul style="list-style-type: none"> • N_p — Number of points per objective pair • n — Number of objective functions <p>Note</p> <ul style="list-style-type: none"> • For problems with two objectives ($n = 2$), $N_{pts} = N_p$ • For problems with three objectives ($n = 3$), $N_{pts} = \frac{N_p(N_p + 1)}{2}$
Shadow minima options and NBI subproblem options	The NBI algorithm uses the <code>fmincon</code> function to solve the shadow minima problem and the NBI subproblems. The options available are similar to those for the <code>fmincon</code> function.

GA Optimization Parameters

The `ga` optimization algorithm in CAGE uses the `ga` function in the Global Optimization Toolbox. In CAGE, `ga` wraps up the `ga` function so that you can use the function for maximizing and minimizing.

Parameter	Description
Generations	Algorithm stops when the number of generations exceeds specified value.
Population size	Number of population members that the algorithm uses. See the Global Optimization Toolbox documentation for guidelines on setting the population size.
Display	<p>Determines the level of diagnostic information displayed in the workspace.</p> <ul style="list-style-type: none"> • (default) <code>none</code> — No information displayed during optimization • <code>iter</code> — Displays statistical information every iteration • <code>final</code> — Displays statistical information at the end of the optimization

Parameter	Description
Crossover function	Function that generates new population members from the existing GA population by crossover. For more information on each function, see the Crossover Options section in the Global Optimization Toolbox documentation. Do not to use a heuristic crossover function for nonlinearly constrained problems.
Crossover fraction	Fraction of the next generation, other than elite children, produced by crossover.
Mutation function	Function that generates new population members from the existing GA population by mutation. The fraction of the next generation, other than elite children, produced by mutation is (1 minus Crossover fraction). For nonlinearly constrained problems, select <code>adaptfeasible</code> .
Selection function	Function that selects the population members that are the parents for the crossover and selection functions.
Hybrid function	Optimization function that runs after the GA termination to try to improve the value of the objective function. If the algorithm has nonlinear constraints, the hybrid function cannot be <code>fminunc</code> or <code>fminsearch</code> . If either of these algorithms are specified for nonlinear constraints, the hybrid algorithm switches to <code>fmincon</code> .
Stall generations	Algorithm stops when the weighted average change in the objective function over Stall generations is less than the Function tolerance .
Stall time limit	Algorithm stops if there is no improvement in the objective function during the specified time stall time limit, in sec.
Function tolerance	Algorithm runs until the weighted average change in the fitness function value over Stall generations is less than Function tolerance .
Constraint tolerance	Tolerance that determines whether a population member is feasible with respect to the nonlinear constraints.
Time limit	Time to stop the algorithm.
Soft constraints weighting	Weighting factor that the optimization uses for soft constraints.

Pattern Search Optimization Parameters

The `patternsearch` optimization algorithm in CAGE uses the `patternsearch` function in the Global Optimization Toolbox. In CAGE, `patternsearch` wraps up the `patternsearch` function so that you can use the function for maximizing and minimizing.

Parameter	Description
Display	Determines the level of diagnostic information displayed in the workspace. <ul style="list-style-type: none"> • (default) <code>none</code> — No information displayed during optimization • <code>iter</code> — Displays statistical information every iteration • <code>final</code> — Displays statistical information at the end of the optimization

Parameter	Description
Time limit	Time to stop the algorithm.
Maximum number of iterations	Maximum number of algorithm iterations.
Maximum function evaluations	Algorithm stops if the number of function evaluations reaches this value.
Variable tolerance	Algorithm stops if the distance between two consecutive free variable values is less than the variable tolerance.
Function tolerance	Algorithm stops if the distance between two consecutive objective function values and the mesh size are both less than Function tolerance .
Constraint tolerance	Determine feasibility with respect to the nonlinear constraints.
Mesh tolerance	Algorithm stops if the mesh size is smaller than this value.
Initial mesh size	Sets the initial size of the mesh for the pattern search algorithm. Do not set this value too small, as insufficient size might lead to the algorithm getting trapped in local optima.
Poll method	Sets the polling strategy that used by the pattern search algorithm. Generally, the <code>GPSPositiveBasis2N</code> and <code>MADSPositiveBasis2N</code> methods are slower than the <code>GPSPositiveBasisNp1</code> and <code>MADSPositiveBasisNp1</code> methods. However, the former methods perform a more thorough search. For more information on these methods, consult the Global Optimization Toolbox documentation.
Search method	Select the function that performs a search, in addition to that performed by the pattern search algorithm. For automotive problems, <code>searchlhs</code> tends to perform well. For details on possible search methods, consult the Global Optimization Toolbox documentation.
Soft constraints weighting	Weighting factor that the optimization uses for soft constraints.

Modal Optimization Parameters

Use the `Modal` optimization algorithm with a composite model to select the best operating mode for each operating point. The algorithm uses the `fmincon` algorithm to optimize an objective for each operating mode and select the best solution.

`Modal` optimization has the same parameters as `fmincon`, plus two additional parameters:

Parameter	Description
Mode variable index	Specify mode variable.
Objective to determine best mode (index)	(Optional) Choose which objective (if you have multiple) to use to select best mode. The default is 1, so <code>CAGE</code> uses the optimized values of the first objective to select the best mode.

MultiStart Optimization Parameters

The `MultiStart` optimization algorithm in `CAGE` uses the `MultiStart` function in the

Global Optimization Toolbox. The `MultiStart` algorithm tries to identify multiple optimal solutions for each operating point. You can set a subset of the algorithm options in CAGE.

In CAGE, the `MultiStart` algorithm generates Sobol design points and chooses the best one to start the optimization. You can specify the number of start points and other options in the Optimization Parameters dialog box.

Parameter	Description
Number of start points	Number of start points per operating point (default is 10).
Start point set type	Choose <code>Sobol Set</code> (space-filling start points) or <code>Random</code> (random start points).
Start points to run	Choose <code>all</code> or <code>bounds - ineqs</code> . Use <code>bounds - ineqs</code> to run only feasible start points that meet constraints.
Run start points in parallel	Choose <code>never</code> or <code>always</code> to run each start point in parallel. Ensure the Distribute Runs optimization option is turned off for the start points to run in parallel.
Tolerance for separate objective values	Specify how far apart objective values must be to qualify as separate local optima.
Tolerance for separate solutions	Specify how far apart solution free variable values must be to qualify as separate solutions.
Local optimization solver	Specify <code>fmincon</code> options.

gamultiobj Optimization Parameters

The `gamultiobj` algorithm uses the `gamultiobj` function in the Global Optimization Toolbox.

Use a point optimization to find feasible start points for a sum optimization, then select **Solution > Create Sum Optimization**. CAGE sets a default population size of 200 for the `gamultiobj` sum optimization. If CAGE does not find a feasible solution, try increasing the population size in the Optimization Parameters dialog box. Larger populations increase the chance of finding feasible points, but take longer to compute.

Scale Optimization

The **Optimization** menu contains the option to **Scale Optimization Items** — Select this to toggle scaling on and off. When you select scaling on, objective and constraint evaluations are (approximately) scaled onto the range [-1 1]. With scaling off, when you run the optimization the objective and constraint evaluations return their raw numbers.

Try running your optimization with scaling off, which is the default setting, to see if it converges to a satisfactory solution (check the output flags and the contour view). If your optimization solution is unsatisfactory, check to see if the objective and constraint functions have vastly different scales. In this case, try turning on scaling, because these optimization problems might benefit from objective and constraint evaluations being scaled to a common scale.

The output view always shows the solutions in raw, unscaled values, whether or not you use scaling to evaluate the problem.

See Also

MultiStart | fmincon | ga | gamultiobj | patternsearch

More About

- “About the NBI (Normal Boundary Intersection) Algorithm” on page 7-23
- “GA Optimization Parameters” on page 7-46
- “Set Up Modal Optimizations” on page 7-25
- “Set Up Multiobjective Optimizations” on page 7-22
- “Set Up MultiStart Optimizations” on page 7-28
- “Global Optimization Toolbox”

Optimization Analysis

- “Using Optimization Results” on page 8-2
- “Filling Tables from Optimization Results” on page 8-7
- “Viewing Your Optimization Results” on page 8-14
- “Analyzing Point Optimization Output” on page 8-27
- “Tools for Optimizations with Multiple Solutions” on page 8-38
- “Analyzing Modal Optimization Results” on page 8-42
- “Analyzing MultiStart Optimization Results” on page 8-48
- “Analyzing Multiobjective Optimization Results” on page 8-51
- “Interpreting Sum Optimization Output” on page 8-54
- “Weighted Objective Pareto Slice” on page 8-60

Using Optimization Results

After you run an optimization, an Output node appears in the optimization tree and the **Optimization Output** views appear. CAGE provides tools for analyzing your results with these views.





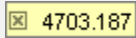
Choosing Acceptable Solutions

CAGE automatically selects successful optimization solutions and highlights unsuccessful solutions for you to investigate. CAGE shows these selections in the icons and check boxes next to the Run column in the Optimization Results table. You can change the selections using the check boxes for each solution, or right-click to change acceptable status of solutions in the graphical views.

This table summarizes how you can use the selected solutions.

Use	See
Filling tables with optimization results.	"Filling Tables from Optimization Results" on page 8-7
Exporting solutions to a data set.	"Exporting to a Data Set" on page 8-3
Importing solutions to other optimization starting values.	"Import from Output" on page 7-32

CAGE indicates the accept status.

Status	Description
	CAGE automatically selects the Accept check boxes for solutions where the algorithm exit flag indicates success (>0). These solutions show a green square icon next to the check box. Typically constraints are met within tolerance.
	Solutions with a red round icon indicate that the algorithm exit flag does not report success (<0). Some constraints might not be met.
	Solutions with an orange triangular icon indicate that the algorithm exit flag is zero. Some constraints might not be met. An exit flag of zero indicates that the algorithm failed because it exceeded computational limits (for example, the algorithm ran out of iterations or function evaluations). You could decide to accept these solutions or you could try changing tolerances and optimizing again.
	Solutions where you have altered the check box status show an asterisk.
	CAGE shows the violated constraints as yellow cells with cross icons in the table. You can control the value used for this highlighting by selecting View > Edit Constraint Tolerance . It is possible to have highlighted constraints within green accept status solutions. The algorithm can report success if constraints are met within tolerance on scaled values. The constraint display applies a tolerance to raw values, and you can also edit this tolerance to help you analyze results. If you are viewing constraints with multiple values and have the view set to Compact, the cell is yellow if any of the individual values are infeasible.

You can view the algorithm output flag in a tooltip by hovering the mouse over each colored accept status icon, or click to select a solution and then you can view the algorithm Exit flag, Exit message, and other details in the Solution Information table.

Create Sum Optimization from Point Optimization Output

Many users employ a point optimization to find good initial values for a sum optimization. To make this workflow easier and faster, you can use a utility to create a sum optimization from your point optimization output.

From your point optimization output node, select **Solution > Create Sum Optimization**.

CAGE creates an optimization (called *Sum_myOptimizationName*). The optimization has these characteristics:

- The objective matches your original optimization but converted to a sum objective.
- The new optimization has identical constraints to your original optimization. Edit or add to these as usual if desired.
- Your original fixed and free variables are converted to a sum optimization (a single run with multiple values).
 - The new optimization uses only accepted solutions from your original optimization output (all runs with a selected **Accept** check box).
 - Therefore, the number of accepted solutions you had in the original optimization determines the number of values within the sum optimization run.
- The free variable initial values and fixed variable values are populated from your point optimal results (accepted solutions only).
- The fixed variables have a **Weights** column with every value set to 1.

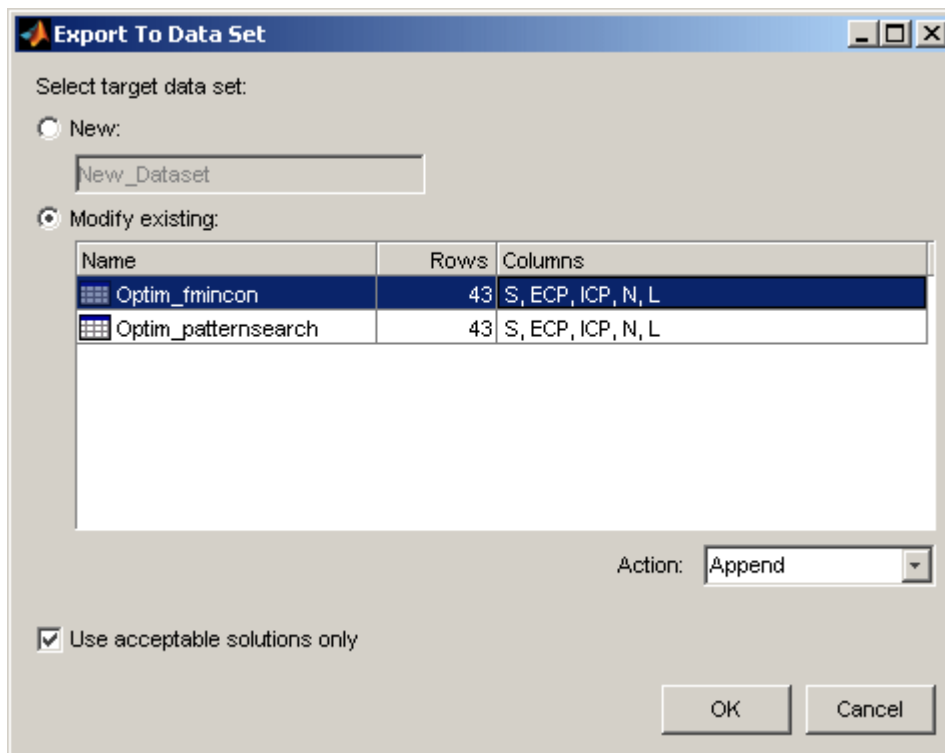
If you rerun point optimizations that have associated sum optimizations, the Model-Based Calibration Toolbox can update and rerun the sum optimization with the point optimization results. The Model-Based Calibration Toolbox uses the existing sum optimization constraints when it reruns the optimization. Previously, the Model-Based Calibration Toolbox created a sum optimization.

Exporting to a Data Set

You can export the optimization output results to new or existing data sets.

To export to a data set:

- 1 Select **Solution > Export to Data Set** or use the toolbar button. The Export to Data Set dialog box appears.



- 2 If exporting to a **New** data set (the default), you can edit the name in the edit box.
- 3 If you want to overwrite or add to an existing data set:
 - a Click the option button **Modify existing**.
 - b Select the desired data set in the list.
 - c Choose from **Action** list:
 - Append adds the data to the chosen data set.
 - Overwrite replaces all data in the data set with the new data.
- 4 By default, the check box **Use acceptable solutions only** is selected. Optimization results with selected **Accept** check boxes are exported. Clear the **Use acceptable solutions only** check box if you want to export all the optimization results.
- 5 Click **OK** and the data is exported to the data set.

This table summarizes the export rules.

Export rules	All fixed and free variables are exported where possible. No models are exported to the data set. If you want to evaluate a model at the variable values, add the model to the data set in the Data Sets view.
--------------	---

When appending, the rules are the same as when merging data sets:

- Columns of inputs are appended to columns with names that match in the data set you are appending to.
- Outputs (models) and any other columns without matching names are not appended.
- The values for any unmatched columns in the data are set to the set point if possible, or zero otherwise.

Custom Fill Function Structure

It is useful to create your own function to fill lookup tables from the results of an optimization, for example, to implement alternative fill methods, smoothing strategies, or to customize output.

The input/output structure of a custom fill function resembles that of the MATLAB interpolation routines INTERP1 and INTERP2. To see the structure of the function, it is best to look at an example:

- 1 Locate and open the file `griddataTableFill.m` in the `mbctraining` folder.
- 2 Type this at the command line to open the example:

```
edit griddataTableFill
```

All 2-D custom fill functions must take these six inputs:

Input	Description
<code>col</code>	Column coordinate of optimization results (NF-by-1)
<code>row</code>	Row coordinate of optimization results (NF-by-1)
<code>fillldata</code>	Optimized results at (row, col) points (NF-by-1)
<code>colaxis</code>	Column breakpoints of table to be filled (1-by-NCOL)
<code>rowaxis</code>	Row breakpoints of table to be filled (NROW-by-1)
<code>currtabdata</code>	Existing table values of table to be filled (NROW-by-NCOL)

The function must pass three output arguments back to CAGE, to allow CAGE to fill the table:

Output	Description
<code>ok</code>	Boolean flag to indicate success of the table fill (TRUE or FALSE)
<code>tabval</code>	New table values of table to be filled (NROW-by-NCOL)
<code>fillmask</code>	Logical matrix to indicate cells to be added to the extrapolation mask as a consequence of the table being filled (NROW-by-NCOL)

In the above specifications:

- NF is the number of points from the optimization results that are used to fill your tables.
- NCOL is the number of column breakpoints in the table.
- NROW is the number of row breakpoints in the table.

Your function should handle the cases when the table fill is successful or not. In `griddataTableFill`, this is handled using the try-catch construct around the call to `griddata`. If `griddata` should fail, then the `ok` flag is set to false and the function returns.

Custom Fill Function for 1-D Tables

You can also write custom fill functions to fill 1-D tables. In this case the input and output specifications are as follows:

Input	Description
<code>row</code>	Row coordinate of optimization results (NF-by-1)
<code>filldata</code>	Optimized results at (row, col) points (NF-by-1)
<code>rowaxis</code>	Row breakpoints of table to be filled (NROW-by-1)
<code>currtabdata</code>	Existing table values of table to be filled (NROW-by-1)

Output	Description
<code>ok</code>	Boolean flag to indicate success of the table fill (TRUE or FALSE)
<code>tabval</code>	New table values of table to be filled (NROW-by-1)
<code>fillmask</code>	Logical matrix to indicate cells to be added to the extrapolation mask as a consequence of the table being filled (NROW-by-1)

See Also

More About

- “Choosing Acceptable Solutions” on page 8-2
- “Creating Sum Optimizations from Modal Optimizations” on page 8-45
- “Creating Sum Optimizations from MultiStart Optimizations” on page 8-50
- “Tools for Optimizations with Multiple Solutions” on page 8-38
- “Viewing Your Optimization Results” on page 8-14

Filling Tables from Optimization Results

In this section...

“Lookup Table Filling from Optimization Results Wizard” on page 8-7


“Lookup Table Filling When Optimization Operating Point Inputs Differ from Table Inputs” on page 8-11

“Filling Tables Via Data Sets” on page 8-13

Lookup Table Filling from Optimization Results Wizard

In a single objective optimization, there is only one solution for each operating point, so you can fill lookup tables with your results. In a multiobjective optimization there is more than one solution per point, and you must first select the preferred solutions before you can use the Lookup Table Filling wizard. To collect your preferred solutions you must use the “Selected Solution Slice” on page 8-39, then you can use this wizard to fill lookup tables with the selected solutions. Modal and multistart optimizations also have multiple solutions per point but CAGE automatically selects solutions for you, so you do not have to select solutions before lookup table filling.


In the Optimization output view, you can use the Lookup Table Filling wizard as follows.

- 1 At the *Optimizationname_Output* node, select **Solution > Fill Lookup Tables**, or click the toolbar button .

The Lookup Table Filling wizard appears.

Note If your tables have been filled before, CAGE remembers all your fill settings between optimization runs and saves the settings with the CAGE project.

- 2 Select the tables to fill, and click the button to add them to the list of tables to be filled. Click **Next**.



- 3 Select or change filling factors for the tables.

CAGE automatically populates the filling factors for the tables if you created your tables using the Create Lookup Tables from Model wizard, and left the defaults to add all your new tables to a tradeoff.

View the **Tradeoff** column to see if a table is associated with a tradeoff. CAGE does not populate the fill factor if a table belongs to more than one tradeoff and there are different fill factors.


Note If you create your tables using the Create Lookup Tables from Model wizard, you can add all your new tables to a tradeoff. The tradeoff can be useful for specifying fill factors for tables, and for investigating optimization results. See “Create Lookup Tables from a Model” on page 3-2.

If your tables are not in a tradeoff and you have not filled them before, select filling factors for your tables as follows:

- a Select a **CAGE table** to be filled in the **Table values to be filled** list.
- b Select the correct variable or model output from the list of **Optimization Results** and click the button to match the result to the table.

Your selected filling factor appears in the **Fill with** column.

Repeat for other tables.




Verify the **Normalizer inputs** show the correct inputs. You might need to specify the **Normalizer inputs** to match with optimization results, if you are filling tables with different inputs to your optimization operating points. See “Lookup Table Filling When Optimization Operating Point Inputs Differ from Table Inputs” on page 8-11.

Click **Next**.

4 Select a **Fill Method**.

- **Extrapolate Fill** — Uses the optimization results to fill the whole table by extrapolation.
- **Direct Fill** — Fills only those table cells whose breakpoints exactly match the optimization points.
- **Clip Fill (column-based)** — Fills table with direct optimization results. Clips the boundary results along the table column.
- **Clip Fill (row-based)** — Fills table with direct optimization results. Clips the boundary results along the table row.
- **Custom Fill** — You can write your own table filling algorithm and use the file browser to select it. See “Custom Fill Function Structure” on page 8-5.



- 5 **Use acceptable solutions only** — Leave this check box selected to use only optimization results marked as 'acceptable'. See “Choosing Acceptable Solutions” on page 8-2.
- 6 **Update tradeoffs** — Select this check box to update tradeoffs with the optimal values from your optimization. You must update your tradeoff to populate it with optimization results. If you do not update the tradeoff, table values and tradeoff values do not match.

For best results, you need a table for each model input (free and fixed, except normalizer variables) to fill simultaneously from the optimization results. You can automatically create a tradeoff with all these tables by using the Create Tables from Model wizard. Evaluation of models in tradeoff uses the variable set points for any variables that do not have a tradeoff table.

- 7 Use the two check boxes on the right to incrementally fill lookup tables from the results of multiple optimizations with smooth interpolation through existing table values. CAGE can extrapolate the optimization results to pass smoothly through table masks and locked cells. Use these features when you want to use separate optimizations to fill different regions of a lookup table.
 - **Use locked table values in extrapolation**— When this check box is selected, CAGE smoothly fills the table between fixed table values and optimization results.

If your calibration tables have fixed values for some table cells, use locked cells for the table cells with fixed values. Such cells often appear on the edge of a table.
 - **Use existing extrapolation mask in fill**— When this check box is selected, CAGE smoothly fills the table between the values in the mask (from previous lookup table filling) and the current optimization results.

Select this check box when you want to fill the same table from multiple optimizations that provide solutions at different operating points. Complex calibration problems can require different optimizations for different regions of a table. The toolbox automatically adds filled cells to the table mask.

If you use the wizard to repeatedly fill a table, CAGE adds to any existing extrapolation mask. As an example, consider filling multiple zones of a table using results from different

optimizations. All zones are cumulatively added to the mask. If there is overlap with previous fills, cells are overwritten unless they are locked. Locked cells are never altered by lookup table filling.


- 8 (Optional) Specify **Filter Rules** to select part of the optimization results for lookup table filling. Specify a filter rule with a logical expression using any input or model available for use in lookup table filling. You can specify an operating mode (for modal optimizations) or any valid expression as a filter. For an example, see “Filling Tables for Operating Modes” on page 8-45.
- 9 Click **Finish** to fill the tables.

A dialog box shows which tables have been successfully filled. Switch to the **Tables** view to examine the tables.

Lookup Table Filling When Optimization Operating Point Inputs Differ from Table Inputs

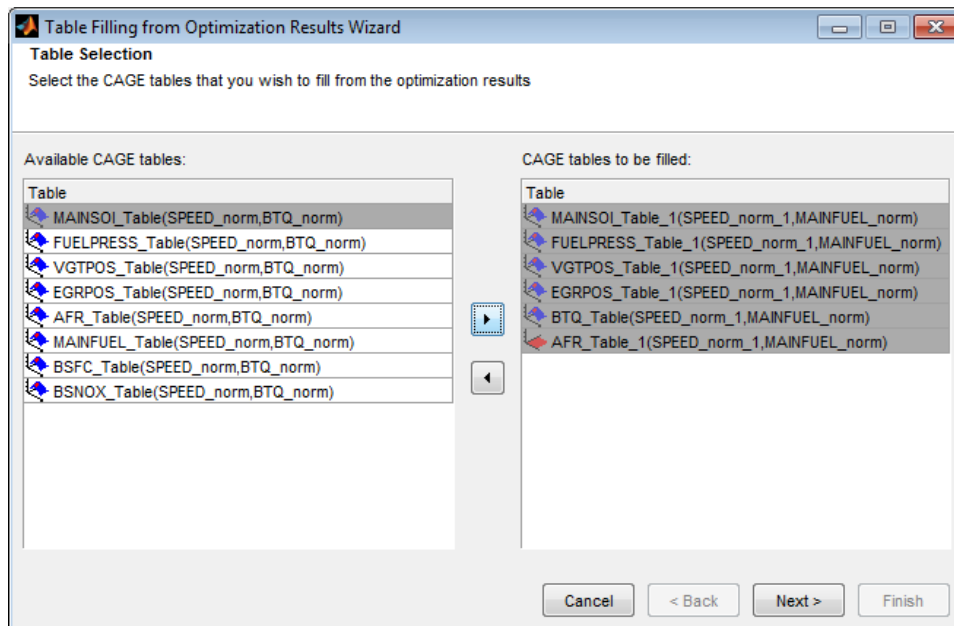
For some optimization problems, you want to optimize at operating points in different variables to the tables you want to fill, and use response models as normalizer inputs to tables. For example, your problem requires running an optimization at torque and speed operating points, but you want to fill lookup tables on axes of mainfuel (a response model) and speed. If all the response model input variables are in your optimization, you can fill tables with that response model as a normalizer input.

- 1 Create your tables using the Create Lookup Tables from Model wizard, and select a response model as an input to your tables. For example, using the example project `DieselPointByPoint.cag` in the `mbctraining` folder, create tables from the `MAINFUEL` model, and select `MAINFUEL` as the Y-axis normalizer input, as shown.




You must specify the breakpoints for your model input normalizer. Edit breakpoints by clicking the button after **Table columns**, and enter a number of points and the range to space the breakpoints over. If you do not do this, model inputs are spaced over 0-1, because CAGE cannot determine the range automatically as happens with variables. After you create your tables with a model input, in the Variable Dictionary you can view a new variable named *modelname_input* with the range you specified. CAGE uses this input variable to match to model names when you fill lookup tables from optimization results.

- When you are ready to fill lookup tables with optimization results, open the Lookup Table Filling from Optimization Results Wizard, select your tables with the response model normalizer input, e.g., MAINFUEL_norm, and click **Next**.



- On the Optimal Result Selection screen, CAGE looks for matches by name among the variables and response models in the Optimization Results list. Verify the Normalizer **Input** column shows the input you want. If CAGE cannot find a match, the **Input** column is empty. To select or change an input, select an item in the Optimization Results list and click the button to select the optimization result for normalizer input.




- Click **Next** and **Finish** to fill your tables.

Filling Tables Via Data Sets

The alternative method of filling tables with optimization output uses Data Sets. This can be useful to see the optimization results and the filled table surface on the same plot. In Data Sets you can also manually edit the results before filling, and compare results with external data.

- From the optimization Output node, click  (Export to Data Set) in the toolbar (or select **Solution > Export to Data Set**). The Export to Data Set dialog box appears. See “Exporting to a Data Set” on page 8-3 for instructions.
- Go to the **Data Sets** view (click the Data Sets button in the **Data Objects** pane) to see that the table of optimization results is contained in the new data set.

You can now use this data set (or any optimization results) to fill lookup tables, as you can with any data set.

- Select the data set and click  (Fill Table from Data Set) in the toolbar.
- Clear the check box to **Show table history after fill**.
- Choose to fill a table with the desired optimization output by selecting them in the two lists, then click the button **Fill Table** at the bottom right.
- Right-click the display and select **Surface** to see the filled table surface and the optimization output values.




See also “Fill Lookup Tables from Data” for an example showing how to use data sets to fill lookup tables.

Viewing Your Optimization Results

Navigating the Optimization Output View

Use the **Optimization Output** view toolbar buttons shown in these figures to determine what is displayed in the table and the graph views. The first default view is the Solution Slice table and the Objective Slice Graphs.

Use these toolbar buttons or the **View** menu to select these Table Views.

Option	See
Toolbar buttons or the View menu. 	<ul style="list-style-type: none"> • “Solution Slice: Optimization Results Table” on page 8-17 and “Solution Slice: Results Surface and Results Contour Views” on page 8-14 • “Pareto Slice Table View” on page 8-38 • “Selected Solution Slice” on page 8-39
Toolbar buttons. You can split and add these views. Use the right-click context menu, the View menu, or the buttons in the view title bars. 	<ul style="list-style-type: none"> • “Objective Slice Graphs” on page 8-18 • “Objective Contour Plot” on page 8-19 • “Constraint Slice Graphs” on page 8-20 • “Constraint Summary Table” on page 8-22 • Free Variable Values Table — Displays the values of the free variables for the currently selected solution. • Solution Information Table — Displays information about the currently selected solution, including the Accept status, the algorithm exit flag and exit message, and other algorithm details such as the number of iterations. <p>Hover the mouse pointer over the Exit message to see the whole message. This message can tell you, for example, if an <code>fmincon</code> optimization run terminated because no feasible start point was found.</p>
Solution menu. 	<ul style="list-style-type: none"> • Select solution — Use this option for choosing your preferred solution for each operating point. • Export to data set — Exports the table visible in the current view only to a new or existing data set. • Fill lookup tables using optimal solutions — Opens the Table Filling From Optimization Results Wizard.

Solution Slice: Results Surface and Results Contour Views

Contour View of Optimization Results






The **Results Contour** view shows a contour plot of one solution at all operating points and all runs in the problem. Use the axes pop-up controls to change what is plotted on each axis. You can plot these against each other:


- Fixed variable settings

- Optimal free variable settings
- Evaluation of objectives at the optimal free variable settings

The optimization results are plotted as points in the contour plot and extrapolation contours (of the z-axis quantity as a function of the x and y-axis quantities) are also displayed.

Each optimization result is displayed using the **Accept** icon, as shown in the Optimization Results table:

-  Successful result
-  Failed result
-  Problem result
-  User-altered accept status.
-  Currently selected result (black outline). Select results by clicking an icon in the plot or a value in the table. Changing the currently selected result in the Results Contour view also updates the result selected in the table, and updates any plots displayed in the lower half of the output view.



Rotation is not permitted in the contour view.

Use the right-click context menu to control these options.

Results to Display	<ul style="list-style-type: none"> • All - Show all optimization results for this solution • Acceptable - Show only the acceptable results for this solution • Green - Show the results with a positive exit flag • Orange - Show the results with a zero exit flag • Red - Show the results with a negative exit flag
Set Acceptable	<p>Mark an optimization result as acceptable if it is marked as unacceptable.</p> <p>Any results whose acceptability has been changed are shown as stars in the plot</p>
Set Unacceptable	Mark an optimization result as unacceptable.
Extrapolate All	Toggles extrapolation from acceptable solutions only (default) to using all results for extrapolation.
Contour Options	<ul style="list-style-type: none"> • Label Contour Lines • Fill Contours • Contour Levels — These contour options are identical to those for the objective contour view. • Show Axes Grid — Toggle whether the axes grid is displayed or not. • Hide Contour — Toggle whether the contour is hidden or not. • Display Contour — Toggle whether the contour is displayed or not.


To toggle between contour and surface view, right-click the view and select **Current View**.

Surface View of Optimization Results

The **Results Surface** view shows a 3D plot of one solution at all operating points and all runs in the problem. Use the axes pop-up controls to change what is plotted on each axis. You can plot these against each other:

- Fixed variable settings
- Optimal free variable settings
- Evaluation of objectives at the optimal free variable settings

The optimization results are plotted as points, and an extrapolation surface (of the z-axis quantity as a function of the x and y-axis quantities) is also displayed. The accept icon for each result is plotted as for the Results Contour.




Left-click anywhere except an icon to rotate the plot.

The right-click context menu shares these options with the **Results Contour** view: **Results to Display**, **Set Acceptable/Unacceptable**, and **Extrapolate All**. Some additional items for the surface view include:

Surface Options	<ul style="list-style-type: none"> • Hide Surface — Toggle whether the surface is visible or not.
------------------------	---

Solution Slice: Optimization Results Table

The Solution Slice view (click ) shows a table with one solution at all operating points and all runs.

This example shows a Solution Slice table display.

Current run: 11

Optimization Results

Vector display format: Expanded vertically

Run	Accept	ICP	N	L	BTQ	BTQ_Bou...
10	<input type="checkbox"/>	40	500	1	214.427	0.895
11	<input checked="" type="checkbox"/>	1.836	1000	0.1	-29.352	0.322
12	<input type="checkbox"/>	3.817	1000	0.2	1.489	0.128
13	<input type="checkbox"/>	17.625	1000	0.3	32.266	0.012
14	<input checked="" type="checkbox"/>	24.668	1000	0.4	60.37	-0.031
15	<input type="checkbox"/>	36.62	1000	0.5	86.937	0.046
16	<input type="checkbox"/>	34.641	1000	0.6	112.713	0.176
17	<input type="checkbox"/>	29.342	1000	0.7	137.876	0.294
18	<input type="checkbox"/>	40	1000	0.8	166.391	0.461
19	<input type="checkbox"/>	40	1000	0.9	192.394	0.609
20	<input type="checkbox"/>	17.581	1000	1	215.279	0.761
21	<input type="checkbox"/>	5.58	1500	0.1	-17.699	0.252
22	<input type="checkbox"/>	5.425	1500	0.2	8.243	9.525e-8
23	<input checked="" type="checkbox"/>	20.168	1500	0.3	35.094	-0.085
24	<input checked="" type="checkbox"/>	39.369	1500	0.4	62.778	3.712e-8
25	<input checked="" type="checkbox"/>	42.607	1500	0.5	90.336	2.115e-7
26	<input type="checkbox"/>	38.436	1500	0.6	115.179	0.011
27	<input type="checkbox"/>	34.312	1500	0.7	140.302	0.161
28	<input type="checkbox"/>	26.331	1500	0.8	164.523	0.324
29	<input type="checkbox"/>	15.128	1500	0.9	188.281	0.472


CAGE automatically selects successful optimization solutions and highlights unsuccessful solutions for you to investigate. CAGE indicates the selections in the icons and check boxes next to the Run column in the Optimization Results table.


The **Solution Slice** view shows a table of one solution at all operating points and all runs in the problem. For single-objective optimizations, there is only one solution per operating point, so the Solution Slice is the only useful view.

For optimizations with more than one solution per run (multiobjective and modal), the solution slice displays controls so you can scroll through the solutions using the arrows or edit box.

The table shows the selected solution at all operating points. The Optimization Results pane shows the fixed variable settings, the optimal free variable settings, and the evaluation of objectives and constraints at the optimal free variable settings.

Objective Slice Graphs

The objective slice graphs are displayed by default for optimization output views, or you can select  in the toolbar.




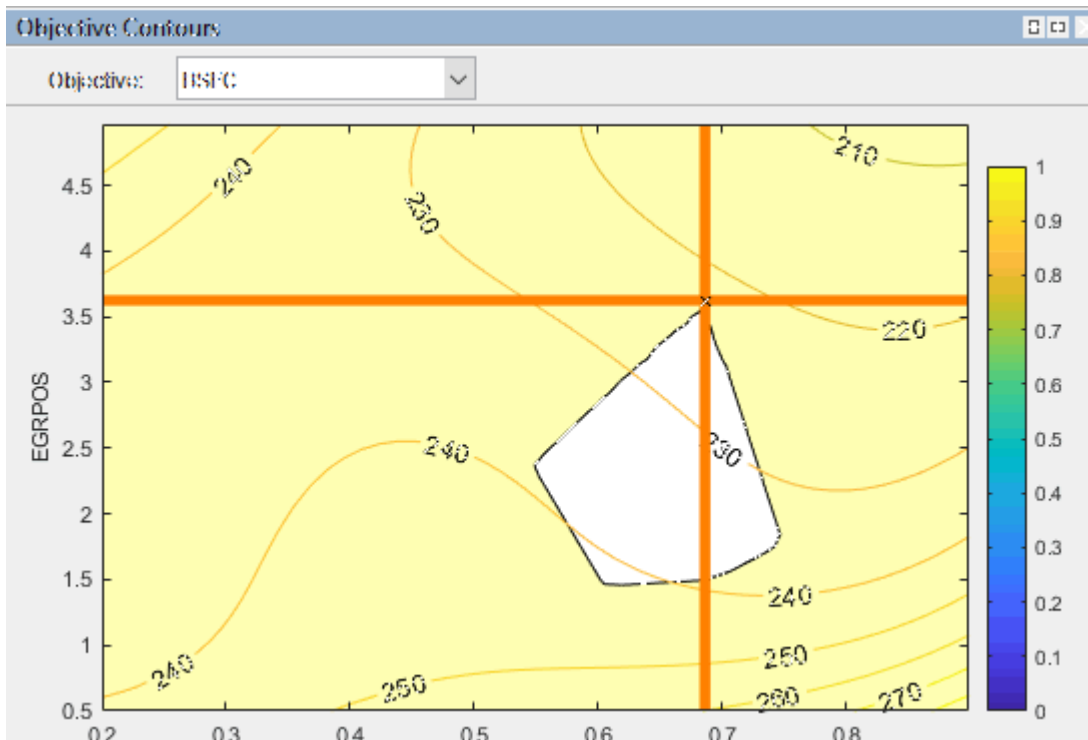
The objective slice graphs show the objective functions at the point selected in the table, with the solution value in orange. Whether the table is displaying a solution slice or Pareto slice, the cell you select in the table is always displayed in the graphs. The objective graphs show cross section plots of the objective function against each free variable in the problem.

The yellow areas show a region outside a constraint tolerance (such as a boundary constraint exported from the Model Browser, or any other optimization constraint). All constraint regions in optimization displays (as in the rest of the toolbox) are shown in yellow.

Use the right-click context menu to toggle constraint display and alter graph size.

Objective Contour Plot


The Objective Contour Plot (click ) shows the contours of the objective against any pair of control parameters, at the run selected in the table, with the solution value at the center of the orange cross-hairs. Yellow areas show a region outside a constraint tolerance. This view is useful for exploring objective functions—a visual way to help avoid local minima.

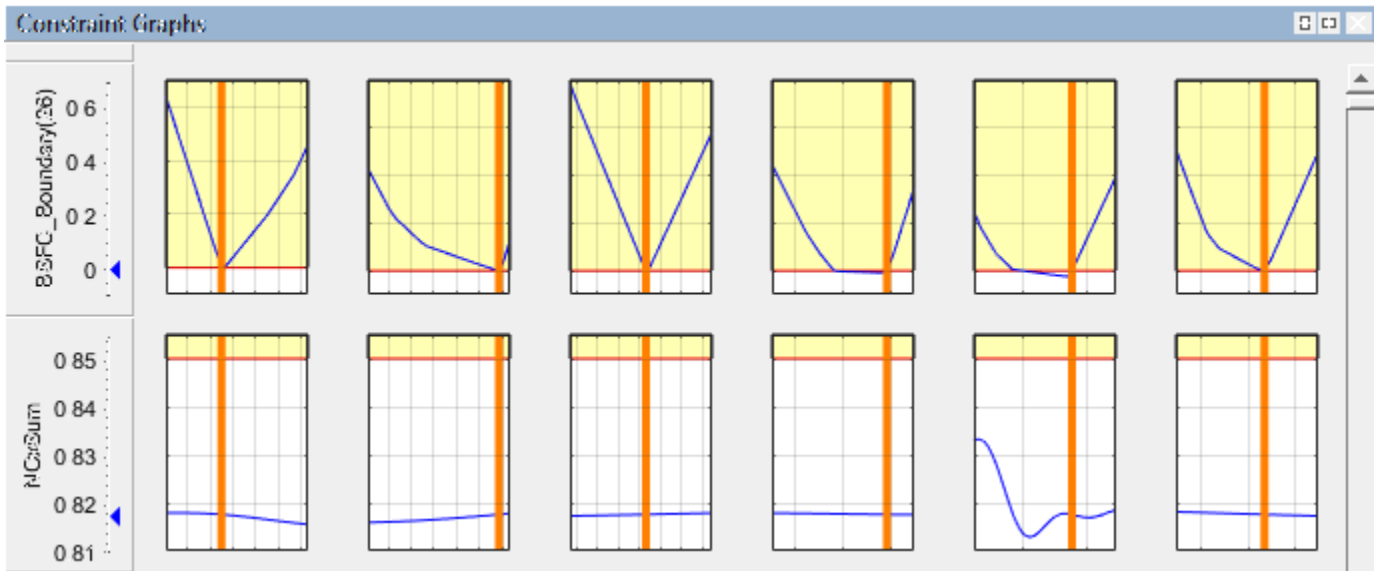


Select parameters to plot in the drop-down lists, and if you have more than one objective you can select from the **Objective** drop-down list.

Use the right-click context menu to toggle constraint display, contour labels, fill contours, and color bar, and control other options.

Constraint Slice Graphs

The Constraint Slice graphs (click ) show the constraint functions at the selected operating point with the solution value in orange. Click inside the tables to select solutions to display. Yellow areas on the graphs show a region outside a constraint tolerance, as shown in this figure.




This example shows the constraint $\text{EXTEMP} \leq 1290^\circ \text{C}$.

The constraint graphs (blue lines) show how the Left Value of each output of a constraint (EXTEMP model) depends on the free variables in the optimization (S, EXH and INT). The Left Value is compared with a plot of the Right Value output (1290°C) on the same axes.


The red horizontal line indicates the Right Value (that is, the upper bound on EXTEMP) which is 1290°C . Because this value is an upper bound, the yellow region above the red line shows where the constraint is infeasible. Yellow is shown above the Right Value plus the tolerance — on many graphs, the distance is too small to see between the red line and the tolerance line. By default, this tolerance is taken from the optimization constraint tolerance. You can control the value used for this highlighting by selecting **View > Edit Constraint Tolerance**.

The vertical orange lines show the optimal values of the free variables. The intersection of these with the blue lines is marked with a blue triangle on the **Constraint1** axis—this intersection is the Left Value (1290°C) at the optimal settings.

If a constraint is violated at the solution value, the Y axis is highlighted in yellow, as shown in Constraint 2 in this example. If constraint values are greater than the tolerance, the row is highlighted in yellow. By default, this tolerance is taken from the optimization constraint tolerance. You can control the value used for this highlighting by selecting **View > Edit Constraint Tolerance**.



Constraint Summary Table

The Constraint Summary Table (click ) view displays the constraint values for the selected solution in the table. This view is useful to see at a glance if a solution met all the constraints. If there are many constraints, it is time-consuming to use the constraint graphs for verification. If you are using equality constraints or tight table gradient constraints, the graphs can appear entirely yellow. You can only see a feasible solution by looking at the Constraint Summary Table, shown in this figure.

Name	Description	Constraint Value	Left Value	Right Value
RESIDFRACatMBT	RESIDFRACatMBT(N, L, ICP, ECP) <= 25	-3.348	21.652	25
MBT_Boundary	Boundary constraint of MBT(N, L, ICP, ECP)	0.454	0.454	0

Constraint values greater than the tolerance appear in bold, and the row is highlighted in yellow. By default, this tolerance is taken from the optimization constraint tolerance. You can control the value used for this highlighting by selecting **View > Edit Constraint Tolerance**. These results should be checked as they might show that the optimization failed to find a solution within the constraint, or they might be within tolerance (very close to zero). Constraint values less than zero are within the constraint.

Constraints are evaluated as inequalities, e.g., the first constraint, RESIDFRACatMBT, as shown in the preceding figure, is $\text{RESIDFRACatMBT} \leq 25\%$. The Left Value shows the left side of the inequality at the optimal settings of the free variables (output of the residual fraction model (RESIDFRACatMBT), which is 21.652). The Right Value shows the right side of the inequality (upper bound, 25%). The constraint value is the difference between the Left and Right values, and the distance to the constraint edge.

In this case, the second constraint, MBT_Boundary, is violated, so the row is yellow, and the positive Constraint Value is highlighted in bold.

Range Constraint Output

The range constraint output is best explained using an example problem.

Description	Variables
Control parameters or free variables	S, EXH, INT
Fixed variables	N, L
Objective	Maximize TQ(S, EXH, INT, N, L) at the fixed values shown in this table.
	Run N L
	1 3000 0.5
	2 4000 0.6
Constraint	Restrict S between an upper and lower bound shown in this table.
	Run N L Min S Max S
	1 3000 0.5 20 30
	2 4000 0.6 30 40

When CAGE runs the optimization, the optimizer returns these optimal values of S, EXH and INT.

Run	N	L	Optimal S	Optimal EXH	Optimal INT
1	3000	0.5	21.33	8.593	29.839
2	4000	0.6	30	5	7.767

CAGE implements this range constraint: Lower Bound (LB) \leq Expression \leq Upper Bound (UB).

Specifically, CAGE implements these expressions.

Description	Expression
Two upper-bound constraints.	$\begin{bmatrix} RangeConLeft(1) \\ RangeConLeft(2) \end{bmatrix} = \begin{bmatrix} -Expression \\ Expression \end{bmatrix}$ $\leq \begin{bmatrix} -LB \\ UB \end{bmatrix} = \begin{bmatrix} RangeConRight(1) \\ RangeConRight(2) \end{bmatrix}$
Range constraint returns two values at each operating point within a run.	$\begin{bmatrix} RangeConOut(1) \\ RangeConOut(2) \end{bmatrix} = \begin{bmatrix} -Expression + LB \\ Expression - UB \end{bmatrix}$
Returned range constraints — Distance from the lower and upper bounds	$RangeConOut(1)$ $RangeConOut(2)$
Constraint	$LB(N, L) \leq S \leq UB(N, L)$
CAGE implementation of constraint	$\begin{bmatrix} -S \\ S \end{bmatrix} \leq \begin{bmatrix} -LB(N, L) \\ UB(N, L) \end{bmatrix}$
Two values at each operating point within a run to the optimizer	$\begin{bmatrix} RangeConOut(1) \\ RangeConOut(2) \end{bmatrix} = \begin{bmatrix} -S + LB(N, L) \\ S - UB(N, L) \end{bmatrix}$

Optimization Output Values										
Vector display format: Expanded vertically										
Run	Accept	S	EXH	INT	N	L	S_Lower...	S_UpperB...	Objective1	Constraint1
1	(1) <input checked="" type="checkbox"/>	21.33	8.597	29.832	3000	0.5	20	30	92.467	-1.33
	(2) <input checked="" type="checkbox"/>									-8.67
2	(1) <input checked="" type="checkbox"/>	30	-5	7.767	4000	0.6	30	40	118.265	0
	(2) <input checked="" type="checkbox"/>									-10

The Optimization Results pane shows the fixed variable settings, the optimal free variable settings, and the evaluation of objectives and constraints at the optimal free variable settings. In this example, the output of the range constraint at the optimal free variable settings is shown in the **Constraint1** column. For each operating point in a run, two values are returned from the range constraint.

For the first run:


Description	Result
Optimal S value	21.33°
Distance from lower bound: <i>RangeConOut(1)</i>	-21.33°+20° = -1.33°
Distance from upper bound: <i>RangeConOut(2)</i>	21.33°-30° = -8.67°

Constraint Output


These are the values shown in the **Constraint1** column. Remember that negative constraint values mean that the constraint is feasible. The same values appear in the Constraint Summary Table for the selected run, in the **Constraint Value** column, as shown in this figure.

Name	Description	Constraint Value	Left Value	Right Value
<input checked="" type="checkbox"/> Constraint1	-S <= -S_LowerBound	-1.33	-21.33	-20
	S <= S_UpperBound	-8.67	21.33	30

The **Constraint Value** gives a measure of the distance to the constraint boundary for each constraint output. If the Left Value > Right Value and greater than the tolerance for any of the constraint outputs, the constraint value is bold and the row is highlighted yellow. By default, CAGE takes this tolerance from the optimization constraint tolerance. You can control the value used for this highlighting by selecting **View > Edit Constraint Tolerance**. This means that this constraint distance should be checked to see if the constraint is feasible at that point.



The Objective Graphs show cross-section plots of the objective function against each free variable in the problem. The left plot is a plot of the objective function against S, with EXH and INT at their optimal values, for the second run. The range constraint for the second operating point ($30 \leq S \leq 40$) is seen; the constraint region is white, and all other regions outside the constraint are yellow.



The constraint graphs for a range constraint shows how the Left Value of each output of a range constraint depends on the free variables in the optimization. The Left Value is compared with a plot of the Right Value output on the same axes. This comparison is illustrated for the example problem at the second run, as shown in the top left graph.

Constraint	Description
Constraint1(1)	<p>First Left Value of the range constraint, <i>RangeConLeft(1)</i>, for the first run in the example problem. The top-left graph shows a blue line, which is a plot of <i>RangeConLeft(1)</i> against <i>S</i>, with all other free variables set to their optimal values. The red horizontal line denotes the Right Value (<i>RangeConRight(1)</i>), that is, the upper bound on <i>S</i> which is -20°.</p> <p>Because this value is an upper bound, the yellow region above the red line shows where the table gradient constraint is infeasible. The vertical orange line shows the optimal value of <i>S</i>; the intersection of this line with the blue line is marked with a blue triangle on the Constraint1(1) axis—the triangle marks the Left Value (-21.3°) at the optimal settings. These are the Left and Right values in the Constraint Summary table for Constraint1(1).</p>
Constraint1(2)	<p>Second Left Value of the range constraint, <i>RangeConLeft(2)</i>, for the first run in the example problem. The bottom left graph shows a blue line plot of <i>RangeConLeft(2)</i> against <i>S</i>, with all other free variables set to their optimal values. The horizontal red line denotes the Right Value (<i>RangeConRight(2)</i>), which is 30°.</p> <p>Because this value is an upper bound, the yellow region above the red line denotes where the table gradient constraint is infeasible. The vertical orange line shows the optimal value of <i>S</i>; the intersection of this with the blue line is marked with a blue triangle on the Constraint1(2) axis—the triangle marks the Left Value (21.3°) at the optimal settings. These are the Left and Right values in the Constraint Summary table for Constraint1(2).</p>

See Also

More About

- “Analyzing Point Optimization Output” on page 8-27
- “Choosing Acceptable Solutions” on page 8-2
- “Constraint Summary” on page 8-55
- “Interpreting Sum Optimization Output” on page 8-54
- “Range Constraint Output” on page 8-23

Analyzing Point Optimization Output

In this section...




“Process for Analyzing Optimization Results” on page 8-27
 “Detecting Local Optima” on page 8-29
 “Investigating Early Termination of Optimization” on page 8-31
 “Handling Flat Optima” on page 8-35

Process for Analyzing Optimization Results

This topic describes a process for analyzing the results from single-objective optimizations (e.g., maximizing torque vs. spark, ICP, ECP at an engine operating point, using `fmincon`, `ga` and `patternsearch` algorithms).

For each run of an optimization, the aim is to find the optimal solution. The Optimization Output View provides graphical tools to help you determine whether an optimal solution has been found for a given run. This view provides a table with icons that indicate the status of each optimization run.


Optimization Results Table Icons

Icon	Description
Green square Accept icon 	Indicates success (algorithm exit flag > 0).
Orange triangle Accept icon 	Indicates the optimization terminated early (exit flag = 0). This situation typically occurs when the optimizer has reached some form of time limit. Examples of this include exceeding a number of iterations or function evaluation limit. In such cases, the optimization was in progress but was forced to terminate before the optimal solution had been found.
Red circle Accept icon 	Indicates failure (an exit flag < 0). Typically this occurs due to the problem being over constrained for this run.

The process for analyzing point optimization results comprises these tasks:





- “Analyzing Output for All Runs” on page 8-27
- “Adjusting Settings To Improve Results” on page 8-28

Analyzing Output for All Runs

- 1 Switch to the Optimization Output view for the optimization.
- 2 Analyze all runs with green square Accept icons (). For each run:
 - a Inspect the Objective Graphs.
 - b Inspect Objective Contour plots for as many pairs of free variables as possible. You can configure the optimization output view to display multiple contour plots simultaneously.

Has the solution found a local optimum? Many optimization algorithms are designed to locate local optima (e.g., `fmincon` in CAGE). Check each successful run to ensure that the optimizer has found the best solution possible. See “Detecting Local Optima” on page 8-29 for more information and examples.

Does the optimization appear to have terminated early? In some cases an optimization appears to return sub-optimal results even though the optimizer has returned a positive exit flag. Investigate such cases. See “Investigating Early Termination of Optimization” on page 8-31.

- 3 Repeat steps 2a and 2b to analyze all runs with orange triangle Accept icons ( ) that indicate the optimization terminated early. See “Investigating Early Termination of Optimization” on page 8-31 for more information and examples.
- 4 Repeat steps 2a and 2b to analyze all runs with red circle Accept icons ( ) that indicate failures. These runs have typically failed to meet constraints. Inspect the plots and determine if it is acceptable to relax any of the constraints.


Adjusting Settings To Improve Results

After you investigate your results to identify problems, use these suggestions to try to improve your optimization results:

- 1 If you detect local optima, try running the optimization again to locate the best optimum.
 - Edit the initial condition manually for this optimization operating point and rerun.
 - For point optimizations that use the `fmincon` algorithm, set the **Number of start points** to be greater than 1 and rerun. In this case, CAGE performs the optimization more than once for each run. To save time, you might want to only repeat the offending runs in this way.
 - Use an alternative algorithm on the runs that have found a local optimum. For example, you could try the `ga` or `patternsearch` algorithms in CAGE (if you tried the `fmincon` algorithm first).
- 2 If the optimization terminates early:
 - In cases where the optimizer runs out of iterations/function evaluations/time and the solution returned is feasible, determine whether the solution is acceptable to you.
 - To accept the solution, select the Accept check box on the Optimization Results table.
 - If you reject the solution, rerun the optimization with modified parameter settings. In this case, if `fmincon` or `patternsearch` is being used, it is advisable to start the optimizer from the solution that has just been found.
 - In cases where the optimizer runs out of iterations/function evaluations/time and the solution returned is infeasible, you can try rerunning the optimization from different initial conditions (for `fmincon` or `patternsearch`) or different parameter settings (all algorithms). If this approach does not resolve the problem, determine if any constraint has been violated. Investigate violated constraints, to determine whether they can be relaxed. If they can, rerun the optimization with the relaxed constraints; if not, leave the check box unselected to indicate the solution is unacceptable.
- 3 See also “Handling Flat Optima” on page 8-35.

Detecting Local Optima

This figure shows views for an optimization which has found the optimal solution. The objective is to maximize Torque (Objective1) against spark angle (S), Exhaust valve closing (ECP) and Intake valve opening (ICP).




When you analyze the optimization results, look for results that have located the best optimum against the free variables.

In this case, an individual plot can only show it is highly likely rather than definitely the optimal value because there are more than two free variables. For problems with more than two free variables, the Objective Graphs and Contours views cannot guarantee that an optimal solution has been found because they provide projections of the model.

For further confirmation, you should inspect the Objective Contour view for as many pairs of free variables as you have time to analyze.


This example shows the algorithm has found a local maximum (marked by the orange cross). You can see the global maximum for this optimization in the lower-right corner of the contour plot.

A constraint, such as a table gradient constraint, could cause a local maximum result. This result could be desirable, however, because it may be preferable for table smoothness to find a local maximum with a slight loss of torque compared to the global maximum (in this case, about 1.3 NM of torque (1%) which is within model accuracy).



To inspect contour plots for many pairs of free variables, you can configure the optimization output view to display multiple contour plots simultaneously. Simultaneous display can help locate those runs that have converged to a local optimum.


This figure simultaneously shows contour plots for all pairs of free variables for the gasoline case study.



Investigating Early Termination of Optimization


Inspect the Objective Graphs and Contour Views to check for optimizations that have terminated early. Early termination typically occurs with runs that have warning orange triangle Accept icons, but can also occur when the optimizer has returned a successful green square Accept icon.

This figure shows an optimization run with a warning orange triangle Accept icon that has been forced to terminate because it exceeded the iterations limit.




In this case, the optimizer has almost found the optimal solution for this run. If this optimizer has taken a long time to run, then as this solution is almost optimal it is probably worth marking as acceptable (select the **Accept** box in the Optimization Results table for this run).

This figure shows another example where an optimization terminated early because it exceeded the iterations limit.




In this case, the problem appears to be over constrained because the plots are entirely shaded yellow. You can check the constraint summary table or the output table to identify if constraints are met. Also inspect the constraint summary and constraint graphs.

The constraint graphs for this case are shown in this figure.



These constraint views confirm that **Constraint2** is violated for this run. Therefore, this solution is probably best left as unacceptable. In cases like this, if it is not already marked as unacceptable, clear the **Accept** box in the Optimization Results table for this run.

This figure shows an optimization that appears to have terminated early despite returning a positive exit flag. You can see that the optimizer has not located the maximum. You should investigate cases like this.



There are many reasons why an optimization appears to terminate early. Two common causes and possible resolutions are discussed in this section.

Poor algorithm parameter settings

`fmincon` may not return a local optimum if these parameter values are too high:

- Variable tolerance
- Function tolerance
- Constraint tolerance

In this case try reducing the values of these parameters to improve performance. However, do not reduce these parameter values too low (less than $\sim 10^{-10}$) to avoid internal issues with `fmincon`. Models that have nonphysical nonlinearity can also cause failure.


Some nongradient-based algorithms may not return an optimum solution. An example of this is the genetic algorithm (`ga`) optimization in CAGE. A poor choice of parameters for such algorithms can lead to early termination of the optimization. For example, setting the **Crossover Fraction** parameter of the `ga` algorithm to 1 can lead to a situation where the algorithm prematurely converges. In this case, try rerunning the optimization at alternative parameter settings. For best results, rerun the algorithm with a Crossover Fraction lower than 1 (the default is 0.8).

Using `fmincon` with noisy models

Optimizations can terminate early because the models are noisy and you used a gradient based algorithm (`fmincon`) to solve the optimization problem.

If the contour plots or any results are suspicious you should always investigate model trends to check if they are sensible and not overfitting. Examine models in the CAGE Surface Viewer or the Model Browser response surface view. You may need to remodel.


To check whether your model is noisy, zoom in on a line plot of the model in the CAGE Surface viewer. Following is a plot of Objective1 against x around the value of x returned by the optimizer.



You can see that the model is noisy and the optimizer has (correctly) returned a local maximum of the model. However, this result is a maximum of the noise component in the model and not the physical component. If the noise is not behavior of the physical system, then you should remodel the noisy models in the Model Browser. The CAGE Import tool can be used to replace the noisy models with the results of the remodeling and the optimization can be rerun.


Handling Flat Optima

Functions that are flat in the vicinity of their optima can be difficult to optimize. This figure shows an example of such a function, $g(x, y) = (x^2 + y^2 + xy)^4$, and its surface plot.



This function has a global minimum at $(0, 0)$ and is very flat in the vicinity of the optimal solution.

Using the `fmincon` algorithm in CAGE to find the minimum of this function (from initial conditions of $(x, y) = [0.5, 0.5]$) produces the result shown in this figure. The optimizer finds a solution at $(x, y) = [-0.113, -0.113]$, which is not optimal. In these plots, you can clearly see that the optimizer has not located the minimum at $(0, 0)$.



To adjust the optimizer to find the minimum, you can take one of several approaches:

- Change the initial conditions.

For a gradient-based algorithm (`fmincon` in CAGE), changing the initial conditions can help the optimizer locate a minimum where the objective function is flat in the vicinity of the minimum. In the example shown in the previous figure, changing the initial conditions to $(x,y) = (1,1)$ leads to `fmincon` finding the minimum at $(0, 0)$.

- Rescale the objective function.

Rescale the objective function with an operation that does not change the location of any optimal solutions, e.g., try taking a square root, fourth root or log, or multiplying by a positive scalar. Check that the position of the optimum is not changed. When an objective function is flat in the vicinity of an optimum, rescaling the objective function can help gradient-based optimization algorithms such as `fmincon` in CAGE. In the example shown in the previous figure, when `fmincon` in CAGE is used to minimize $10^{12}g(x, y)$, the minimum at $(0, 0)$ is located.

- Use a non-gradient based algorithm.

Try either the pattern search or genetic algorithm options. As these algorithms do not use gradient information, they can perform better when used on optimization problems with flat minima. In the example shown in the previous figure, the pattern search algorithm in CAGE located the minimum using the default settings.

- Run the optimization from several initial condition values.

If you are using `fmincon` then another possible workaround is to set the **Number of Start Points** parameter to be greater than 1. This setting runs `fmincon` the specified number of times from different start conditions. Use this option only for the affected runs as it can be time consuming.

- Change tolerances.

For a gradient-based algorithm (`fmincon` in CAGE), changing the variable or function tolerances can help the optimizer locate a minimum where the objective function is flat in the vicinity of the minimum. Reducing the **variable** and **function** tolerances may improve the convergence to the optimum value in this case.

Tools for Optimizations with Multiple Solutions

In this section...

“Analyzing Modal, MultiStart, and Multiobjective Optimizations” on page 8-38

“Pareto Slice Table View” on page 8-38

“Selected Solution Slice” on page 8-39

“Exporting Selected Solutions” on page 8-41

Analyzing Modal, MultiStart, and Multiobjective Optimizations


CAGE has additional tools for analyzing optimizations with more than one solution for each operating point. Optimizations with multiple solutions are multiobjective optimizations, modal optimizations and multistart optimizations. Use the optimization output node tools to view all solutions and select solutions. The tools for viewing and selecting solutions are described in these sections:

- “Pareto Slice Table View” on page 8-38 shows a table of all solutions at one run.
- “Selected Solution Slice” on page 8-39 is for collecting and exporting only the solutions you have decided are optimal at each run.

You can export selected solutions or all solutions to a data set, and you can restrict export to acceptable solutions only. See “Exporting Selected Solutions” on page 8-41.


- Check the messages and exit flags for each solution, shown in the Optimization Results table and the Solution Information pane.
- For advice on multiobjective optimizations, see “Analyzing Multiobjective Optimization Results” on page 8-51 .
- For advice on modal optimizations, see “Analyzing Modal Optimization Results” on page 8-42.
- For advice on multistart optimizations, see “Analyzing MultiStart Optimization Results” on page 8-48.


Pareto Slice Table View

The Pareto Slice table view (click ) is for optimizations where there is more than one solution at each run (multiobjective, multistart or modal). The Pareto Slice shows a table of all solutions at one run; you can scroll through the runs using the arrows or edit box at the top.


To collect best solutions across different runs, you need to select a solution for each run, and your selections are stored in the Selected Solution slice.

To select a solution for each run:

- 1 Decide which solution you want to use for the currently selected run. Use these tools to help you:
 - Use the Pareto Slice table and Results contour and surface views along with the “Objective Slice Graphs” on page 8-18 to select the best solution for the run. If you have constraints you can also use the “Constraint Slice Graphs” on page 8-20 and “Constraint Summary Table” on page 8-22 to help you decide which solution to choose for each run.
 - For multiobjective optimizations, display the “Pareto Graphs” on page 8-51 (click  in the toolbar) which shows the available solutions with the current selection highlighted in red.

- For modal optimizations, see “Analyzing Modal Optimization Results” on page 8-42.
- 2 When you have decided which solution you want to use for the currently selected run, you can select it as best by editing the **Selected solution** control above the table, or by clicking Select Solution () in the toolbar. You can also select best solutions with the toolbar in the Solution Slice view, see “Solution Slice: Optimization Results Table” on page 8-17.
 - 3 Scroll through the runs and select a best solution for each. These selections are collected in the Selected Solutions Slice, where you can view them, use them to fill tables, or export to a data set. You can also import them to an optimization. See “Selected Solution Slice” on page 8-39.

Before you run an NBI optimization you can specify how many solutions you want the optimization to find, using the Set Up and Run Optimization toolbar button.

Solution	 Accept	grackmea	egrift	tq_desired	afr_min	measrpm	Objective1	Objective2	Constraint1	Con
1	<input type="checkbox"/>	-16.604	-282.651	200	30	5000	0.39	1.668	1.239e6	-
2	<input type="checkbox"/>	-0.92	-9.753	200	30	5000	-78.305	-1.789e-3	-4815.248	4
3	<input type="checkbox"/>	-0.426	-27.69	200	30	5000	56.569	4.074e-3	4703.187	-4
4	<input type="checkbox"/>	-0.891	-4.196	200	30	5000	-72.481	-1.189e-3	-4029.09	
5	<input checked="" type="checkbox"/>	-0.398	-6.123	200	30	5000	-100.689	-2.259e-4	-2958.292	2
6	<input type="checkbox"/>	-0.648	-1.815	200	30	5000	-94.975	-6.094e-4	-3125.105	3
7	<input type="checkbox"/>	-0.85	-1.644	200	30	5000	-77.764	-7.608e-4	-3773.391	3
8	<input checked="" type="checkbox"/>	-0.811	-1.577	200	30	5000	-81.286	-7.731e-4	-3619.442	3
9	<input type="checkbox"/>	-0.796	-1.552	200	30	5000	-82.665	-7.778e-4	-3563.195	3
10	<input type="checkbox"/>	429.743	949.514	200	30	5000	-1.121e-5	17473.161	5.108e10	-5

As in the other table views, you can use the Accept check boxes to choose a selection of rows within the table. In this table view, you can only use this to select solutions within a single run. Each different solution has a check box and colored icon for “Acceptable” status. You can override these selections using the check boxes if you want to choose solutions within a run, for use when exporting to a data set, importing to other optimizations, or for future reference. See “Choosing Acceptable Solutions” on page 8-2.

Selected Solution Slice

In a multiobjective, modal or multistart optimization, there is more than one possible optimal solution at each run. You can use the **Selected Solution** view to collect, view, and export those solutions you have decided are optimal at each run.



Click Selected Solution in the toolbar  to view the Selected Solution view.

CAGE selects solutions depending on the type of optimization as follows:


- Modal optimizations and MultiStart optimizations select a solution for each run automatically that you can view and change manually if you want.
- For multiobjective optimizations you must choose solutions manually to decide the acceptable tradeoff between the competing objectives.

Note You can use the Selected Solution slice, or alternatively you can choose to export all solutions to a data set. See “Exporting Selected Solutions” on page 8-41

You can use the plots and table views to help you select best solutions for each run. These solutions are saved in the **Selected Solution** view. You can then export your chosen optimization output for each point from the **Selected Solution** view to a data set, or use your optimization output to fill lookup tables or import to another optimization.


- 1 If you want you can initialize the Selected Solutions view with a particular solution for all runs. Select **Solution > Selected Solution > Initialize**.

The **Create Selected Solution** dialog box appears.



The default 1 initializes the first solution for each run as the selected solution. You can edit the solution number here if you want. For example, if you select 4, solution number 4 is initialized as the best solution for every run. Click **OK**.

- 2 Use the table views and the plots in the graphs (Objective Slice, Pareto Front, and Constraint Slice graphs) to help you select the best solution for each run. Use the procedure described in “Pareto Slice Table View” on page 8-38 to select a solution for each run. Repeat until you have selected solutions for all runs.
- 3 You can also change selected solution in the Selected Solution slice view, by editing the **Selected solution** control above the table. You should use the Pareto Slice table and other views to investigate all solutions.

These solutions are saved in the **Selected Solutions** view. This view collects all your selected solutions together in one place. For example, you might want to select solution 7 for the first run, and solution 6 for the second, and so on. You can then use your chosen optimization output for each point to fill lookup tables (see “Filling Tables from Optimization Results” on page 8-7), or choose the Export to Data Set  toolbar and **Solution** menu option (see “Exporting to a Data Set” on page 8-3), or use these solutions as starting points in another optimization (see “Import from Output” on page 7-32).

An example of the **Selected Solutions** view is shown. It looks similar to the Solution Slice view, except the **Selected solution** controls at the top are enabled instead of the **Current solution**. You can change the selected solution in this view. The solution chosen as best (in this or other views) for the currently selected run is displayed in both current and selected solution edit boxes.

As in the other table views you can use the Accept check boxes to choose a selection of rows within the table. See “Choosing Acceptable Solutions” on page 8-2.

Run	Accept	S	ICP	ECP	N	L	BTQ	BTQ Bo...
10	<input type="checkbox"/>	50	23.877	22.429	5504.174	0.092	24.56	0.786
11	<input type="checkbox"/>	33.746	23.014	16.39	879.31	0.178	46.638	6.256e-3
12	<input checked="" type="checkbox"/>	35.653	17.897	22.411	1393.184	0.178	50.103	-0.202
13	<input checked="" type="checkbox"/>	41.293	19.8	27.831	1907.058	0.178	52.258	-0.276
14	<input checked="" type="checkbox"/>	41.938	23.186	26.747	2420.931	0.178	52.832	-0.336
15	<input checked="" type="checkbox"/>	38.544	24.748	18.762	2934.805	0.178	52.86	-0.465
16	<input checked="" type="checkbox"/>	34.641	22.01	13.787	3448.679	0.178	51.998	-0.345
17	<input checked="" type="checkbox"/>	42.135	38.387	2.765	3962.552	0.178	52.058	1.967e-8
18	<input checked="" type="checkbox"/>	46.259	36.71	6.226	4476.426	0.178	51.398	1.053e-9
19	<input type="checkbox"/>	50	23.877	22.429	4990.3	0.178	49.397	0.029
20	<input type="checkbox"/>	50	23.877	22.429	5504.174	0.178	49.252	0.256
21	<input type="checkbox"/>	24.496	23.877	22.429	879.31	0.265	92.897	0.138
22	<input checked="" type="checkbox"/>	25.183	21.542	32.679	1393.184	0.265	96.452	-0.089

Exporting Selected Solutions

For optimizations with multiple solutions (multiobjective, modal and multistart), you can choose to export only selected solutions or all solutions to a data set. You can restrict export to acceptable solutions only (specified by the Accept check boxes). See “Choosing Acceptable Solutions” on page 8-2. You can use the Selected Solution slice to collect only the best solution for each run.

- 1 Select **Solution > Export to Data Set** or use the toolbar button. The Export to Data Set dialog box appears. For optimizations with multiple solutions this dialog provides an additional control called **Solutions to Export**.
- 2 Use the **Solutions to Export** drop-down list to select either:
 - **Selected Solutions** — this exports your collected solutions in the Selected Solution Slice.
 - **All Solutions** — this exports every solution from every run.
 - You can choose whether to export acceptable solutions only with the check box **Use acceptable solutions only**.

For more information on exporting optimization results, see “Exporting to a Data Set” on page 8-3.

See Also

Related Examples

- “Analyzing Modal Optimization Results” on page 8-42
- “Analyzing MultiStart Optimization Results” on page 8-48
- “Analyzing Multiobjective Optimization Results” on page 8-51

Analyzing Modal Optimization Results

After you run your modal optimization, use the optimization output node to verify the results. For general advice, see “Analyzing Point Optimization Output” on page 8-27. This describes analyzing *modal* optimization results.

Viewing and Selecting Modal Optimization Results

Modal optimization results have more than one solution at each operating point. The modal optimization algorithm tries to automatically select the best mode for each operating point.

Use the optimization output node tools to view all solutions, see which solution is selected, and change the selections manually if you want. These features are also useful for selecting solutions for multiple objective optimizations (using the NBI algorithm) and multiple start points (using the `MultiStart` algorithm) that also have more than one solution per point.

- 1 The default view in the `GasolineComposite.cag` example `BTQ_Optimization_Output` optimization shows contour plots for mode 1, the 4 cylinder mode. You can view the Results Contour, Objective Contours, Optimization Results, and Constraint Summary in the layout.

Current run: 1 Current mode: 1 Selected 1 Type: Modal

All Optimization Results

Results Contour

X-axis: N Y-axis: L Z-axis: ...

Current Result - Optimization Solution

Objective Contours

Objective: BTQ

X-axis: S Y-axis: ICP

Optimization Results

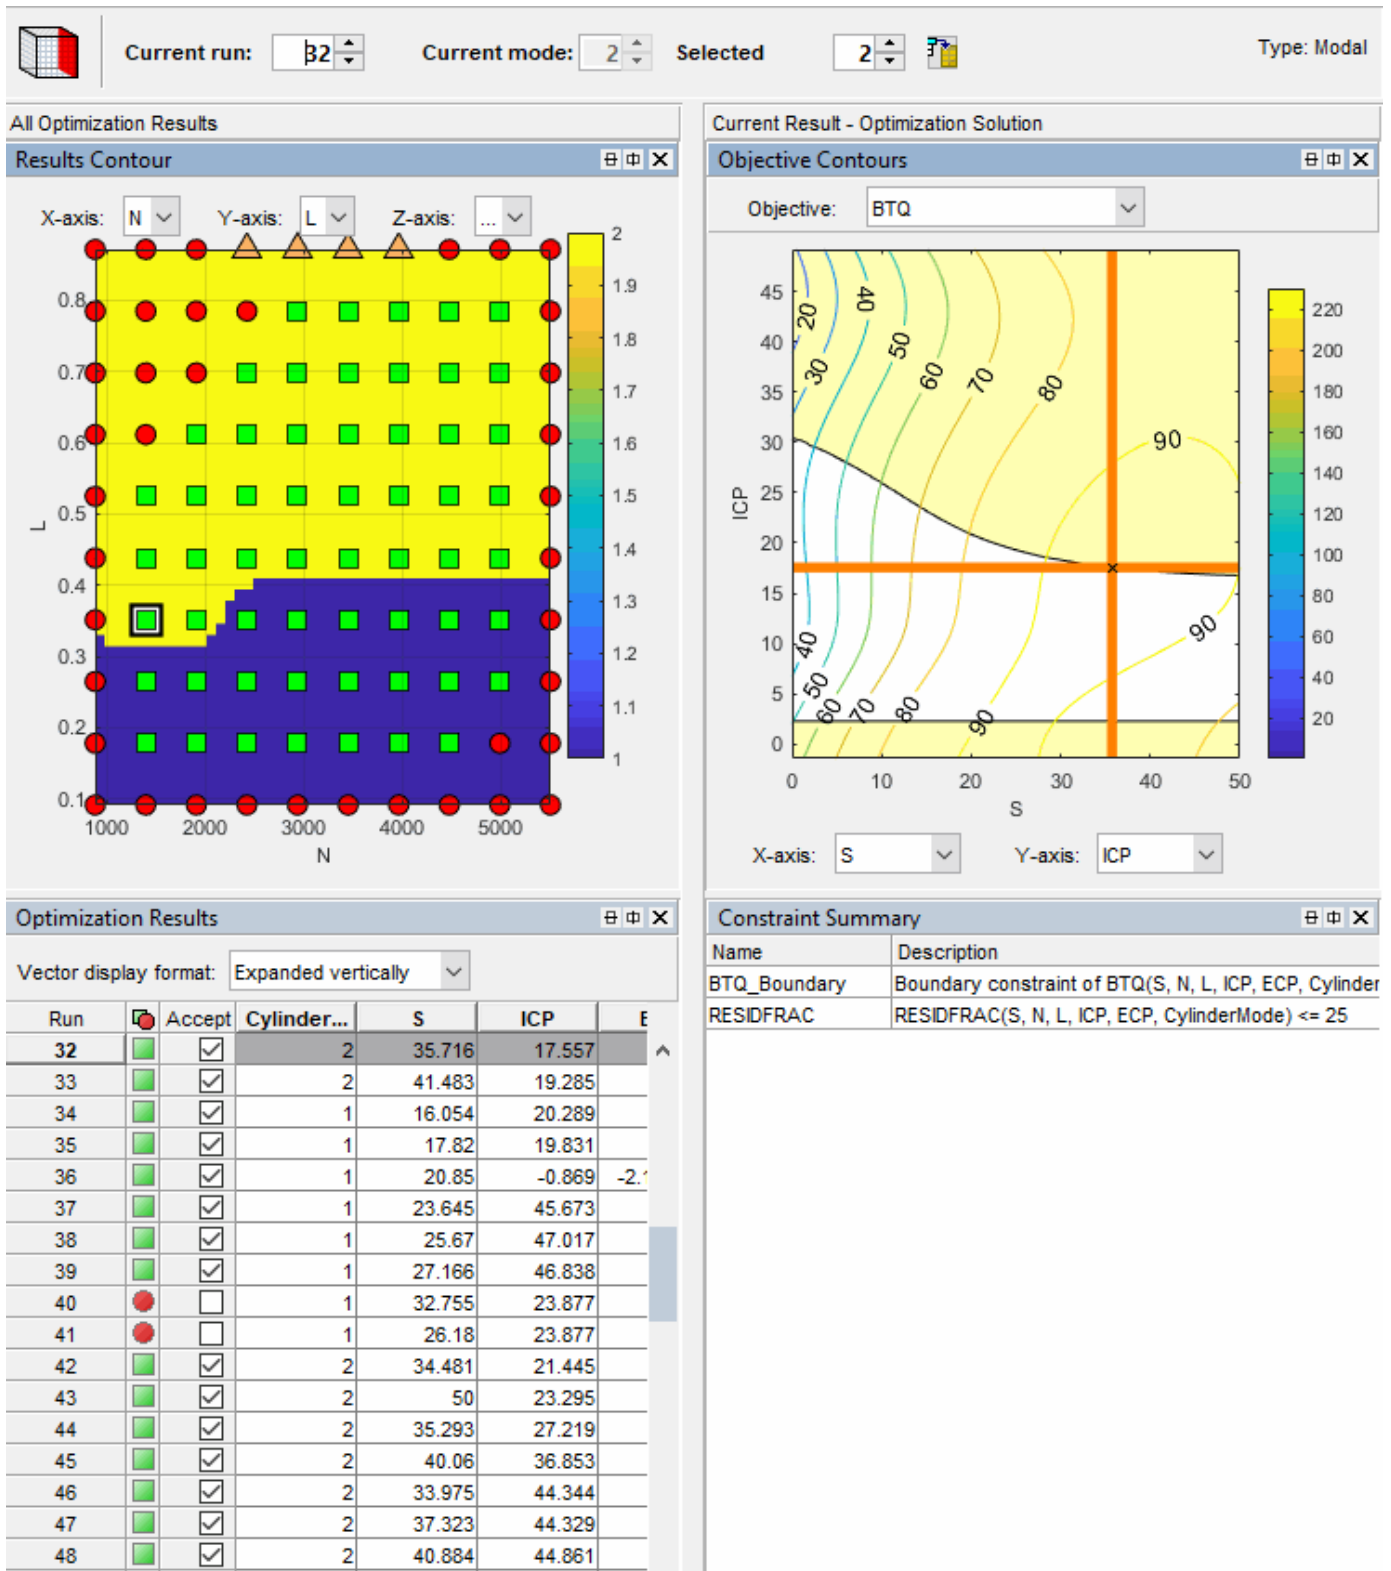
Vector display format: Expanded vertically

Run	Accept	Cylinder...	S	ICP	E
1	<input type="checkbox"/>	1	45.503	2.678	
2	<input type="checkbox"/>	1	48.452	0.536	
3	<input type="checkbox"/>	1	47.986	5.621	
4	<input type="checkbox"/>	1	46.999	21.407	
5	<input type="checkbox"/>	1	47.716	21.876	
6	<input type="checkbox"/>	1	49.767	20.734	
7	<input type="checkbox"/>	1	50	23.877	
8	<input type="checkbox"/>	1	50	23.877	
9	<input type="checkbox"/>	2	50	23.877	
10	<input type="checkbox"/>	2	50	23.877	
11	<input type="checkbox"/>	1	33.746	23.014	
12	<input checked="" type="checkbox"/>	1	35.653	17.897	
13	<input checked="" type="checkbox"/>	1	41.293	19.8	
14	<input checked="" type="checkbox"/>	1	41.938	23.186	
15	<input checked="" type="checkbox"/>	1	38.544	24.748	
16	<input checked="" type="checkbox"/>	1	34.641	22.01	
17	<input checked="" type="checkbox"/>	1	42.135	38.387	
18	<input checked="" type="checkbox"/>	1	46.259	36.71	

Constraint Summary

Name	Description
BTQ_Boundary	Boundary constraint of BTQ(S, N, L, ICP, ECP, Cylinder
RESIDFRAC	RESIDFRAC(S, N, L, ICP, ECP, CylinderMode) <= 25

- Select one of the accepted results. The table and contour plot display the selected best solution for all operating points.



- 3 Review the Results Contour plot to see which mode has been selected across all operating points. Use this view to verify the distribution of mode selection.
- 4 If you have extra objectives, you can also view them in the tables and plots. Use the other objectives to explore the results. For example, you might want to manually change the selected mode based on an extra objective value. If you have extra objectives, it can be useful to view plots of the other objective values at your selected solutions.
- 5 Click to select a point in the table or Results Contour. Use the **Selected solution** controls to alter which mode is selected at that point. You might want to change selected mode if another mode is also feasible at that point. For example, you can change the mode if you want to make the table more smooth.

In the `GasolineComposite.cag` example, you can run some operating points in either 4- or 8-cylinder mode. When both modes are feasible, the modal optimization algorithm selects the mode that results in the best torque.

- 6 Use the Pareto Slice view to see all the solutions for a particular operating point. You can inspect the objective value (and any extra objective values) for each solution. If needed, you can manually change the selected mode to meet other criteria, such as the mode in adjacent operating points, or the value of an extra objective. Change the selected solution using the **Selected solution** control or by selecting the solution and using the toolbar.
- 7 If you change the selected mode for a point, return to the Selected Solution view to observe the selected solutions for all operating points.
- 8 Check the messages and exit flags for each solution, shown in the Optimization Results table (hover over the Accept icons) and the Solution Information pane. Modal optimizations provide exit messages from `fmincon` and prefix the message with the mode number for the solution. See the `fmincon` function for exit messages. There is also an exit message specific to modal optimization: -7 which reports that the mode is not valid (NaN) for a particular operating point.

Creating Sum Optimizations from Modal Optimizations

When you are satisfied with all selected solutions for your modal optimization, you can make a sum optimization over all operating points. The mode must be fixed in the sum optimization to avoid optimizing a very large number of combinations of operating modes. For example, the `GasolineComposite.cag` example optimization has $2 \times 57 = 114$ different combinations of modes.

To create a sum optimization from your point modal optimization:

- 1 From your point optimization output node, select **Solution > Create Sum Optimization**.

The toolbox automatically creates a sum optimization for you with your selected best mode for each operating point. The create sum optimization function converts the modal optimization to a standard single objective optimization (`fmincon` algorithm) and changes the Mode Variable to a fixed variable.

- 2 You can then add table gradient constraints to ensure smooth control and engine response.

Filling Tables for Operating Modes

Composite models can require the ability to select part of the optimization results to fill a particular table. For example, you need to discard solutions for other modes when filling a table with an input that is not used for all modes.

You can apply filter rules to select part of the optimization results for lookup table filling. The filter rules are important for modal optimizations. You can specify an operating mode or any valid expression as a filter when using the Table Filling wizard.

- Use filter rules when your goal is to fill a different table for each mode.
- Specify a filter rule with a logical expression using any input or model available for use in lookup table filling.
- The Lookup Table Filling from Optimization Results wizard automatically sets up filter rules for you if some inputs are not used for all modes in your composite model.

From any type of optimization, you can use the Lookup Table Filling From Optimization Results Wizard. The example project `CompositeWith2Tables.cag` shows the use of filter rules in the wizard to specify results from a single mode to fill a specified table.

In this example project:

- There is a single table for each control variable which stores the value for the best mode. The strategy has separate tables for each mode.

Composite calibration problems of this kind often involve separate optimizations (point and sum) with different free variables and constraints for each mode.


- There is a separate point optimization for each mode. The results from each mode are exported to the same data set (using the append option). The sum optimization uses the point results data set.
- To finish off the calibration, the sum optimization provides results for a multimodal drive cycle, using the selected mode at each point.

To see the example:

- 1 Load the example project `CompositeWith2Tables.cag` found in `matlab\toolbox\mbc\mbctraining`.
- 2 View completed examples of composite models, optimizations, and filled tables.
- 3 To see the lookup table filling filter rules, expand the `Sum_BTQ_Optimization` node to view the optimization output node.
- 4 Select **Solution > Fill Lookup Tables** or use the toolbar button.

The Lookup Table Filling From Optimization Results Wizard appears.

- 5 Click **Next** to review the saved settings in the wizard.
- 6 On the final screen of the wizard, you can view filter rules. These rules specify which mode to use to fill each table.



See Also

More About

- “Analyzing Point Optimization Output” on page 8-27
- “Create Sum Optimization from Point Optimization Output” on page 8-3
- “Filling Tables from Optimization Results” on page 8-7

Analyzing MultiStart Optimization Results

In this section...

“Viewing and Selecting MultiStart Results” on page 8-48


“Creating Sum Optimizations from MultiStart Optimizations” on page 8-50

Viewing and Selecting MultiStart Results

After you run your optimization, use the optimization output node to verify the results. For general advice, see “Analyzing Point Optimization Output” on page 8-27. This process describes features specific to the results of MultiStart optimizations.

Optimizations using the MultiStart algorithm have multiple start points and try to find multiple solutions per point. CAGE selects the best solution based on the objective value. You can investigate all solutions and change selected solutions manually if you want, for example to make smoother tables.

To examine MultiStart optimization results:


- 1 Click the Selected Solution button  in the toolbar to see the optimal results selected by CAGE in the Selected Solution table.
- 2 View your results in the Results Contour plot. Look for table areas that are not smooth enough.

You can also view the Results Surface at the same time by right-clicking the title bar and selecting **Split View Horizontally**.

Focus on runs that have accepted solutions (green squares) and then solutions that ran out of time (orange triangles). Red circles indicate failures to meet constraints with any of the start points (e.g., outside boundary model), so further analysis is less useful compared to the accepted solutions. For example, investigate green squares where the table is not very smooth.

- 3 Click the plots or table to select a point to investigate.

This example shows a selected point where the value of spark is too different from the neighboring points, which makes the table not smooth enough.



- 4 Click  in the toolbar to switch to the Pareto Slice and view all solutions at the selected point.


This example shows MultiStart results as follows:

- CAGE sorts MultiStart results with the best solution at the top (solution 1).
- The number of solutions is not necessarily the same as the **Number of start points**. The example has five feasible solutions, and an additional row displaying NaNs. This means that CAGE found six different feasible solutions for at least one other run in this optimization. Ignore any rows with NaNs. CAGE shows the same maximum number of solution rows for every run. If there are rows beyond the feasible solutions for the current run, then CAGE fills the rows with NaNs.

You can set the tolerance between different solutions with the **Tolerance for separate solutions** MultiStart setting.

Here, CAGE has selected the best solution with the optimal value of torque, BTQ. In this case you can instead select another solution to make a smoother table in spark (S) with only a small tradeoff in the torque value.

Solution	Accept	S	ICP	ECP	N	L	BTQ	BTQ_Bo...	RESIDFR...
1	<input checked="" type="checkbox"/>	40.384	20.168	20.815	1500	0.3	35.094	-0.085	9.113e-7
2	<input checked="" type="checkbox"/>	24.343	3.507	19.167	1500	0.3	34.793	-2.555e-9	-7.745
3	<input checked="" type="checkbox"/>	30.984	26.726	7.708	1500	0.3	34.489	3.528e-9	-4.377
4	<input checked="" type="checkbox"/>	31.073	26.762	7.734	1500	0.3	34.489	3.427e-8	-4.346
5	<input checked="" type="checkbox"/>	1.608e-20	22.363	27.613	1500	0.3	28.788	-0.024	2.541e-10
6	<input type="checkbox"/>	NaN	NaN	NaN	1500	0.3	NaN	-0.145	NaN

- Change the selected solution using the **Selected solution** control, or click the solution in the table and click Select Solution  in the toolbar.
- Return to the Selected Solution slice to view the difference in your table.
- Repeat the process to investigate your other results.

Creating Sum Optimizations from MultiStart Optimizations

When you are satisfied with all selected solutions for your optimization, you can make a sum optimization over all operating points. To create a sum optimization from your point MultiStart optimization:

- From your point optimization output node, select **Solution > Create Sum Optimization**.

The toolbox creates a sum optimization with your selected solution values defining the operating points. The create sum optimization function converts the MultiStart optimization to a standard single objective sum optimization (fmincon algorithm) and uses your accepted selected solutions for variable values.

- Add table gradient constraints to ensure smooth control and engine response.

See also “Create Sum Optimization from Point Optimization Output” on page 8-3.

Analyzing Multiobjective Optimization Results

To analyze the multiobjective optimization results when there is more than one solution for a run, use the Pareto Graphs.


Pareto Graphs

Pareto graphs show the available solutions for the selected run:


- CAGE highlights the current selection with a gray circle outline.
- Pareto solutions are green squares.
- Dominated solutions are orange triangles.
- Unsuccessful solutions are red circles (constraints may not be met or the algorithm ran out of time).
- The selected solution is a star shape.

Click in the tables or graphs to view solutions. To change the selected solution, edit the number in the **Selected solution** box. CAGE displays the selected solution in the other graphs (objective and constraint).

Use the Pareto graphs with the table views (Solution Slice and Pareto Slice) and the other plots in the graphs (Objective Slice and Constraint Slice graphs) to help you select best solutions for each run. You can collect these solutions together in the “Selected Solution Slice” on page 8-39.



Tip It is useful to display the Solution Information view at the same time to view information about a selected solution. You might want to select a dominated solution (orange triangle) over a pareto solution (green square) to trade off desired properties.



Before you run an NBI multiobjective optimization, you can specify how many solutions you want the optimization to find by using the Optimization Parameters dialog box. The NBI default number of solutions is 10 for two objectives. Click **Set Up and Run Optimization** to change the number of solutions before running again. In the Optimization Parameters dialog box, change how many tradeoff solutions you want the optimization to find per run.

For `gamultiobj`, if CAGE does not find a feasible solution, you might try increasing the default population size.

Multiobjective Output Messages

Multiobjective solutions can have specific exit messages.

To view the exit flag, hover over the accept status icon (green square, orange triangle, or red circle) for a run in the Optimization output view. You can view the algorithm exit messages for the currently selected run by selecting **View > Current View > Solution Information**. Check these messages to check for problems with your optimization.

Message	Description
Multiobjective <code>gamultiobj</code> Output	See <code>gamultiobj</code> .

Message	Description												
Multiobjective NBI Output	<p>CAGE displays shadow solutions at the start of the solution list and indicated by the prefix “Shadow solution” in the message.</p> <p>NBI output messages include the exit flags and the first part of the message returned by <code>fmincon</code> calls. CAGE adds extra information about the NBI solution to the message. The extra NBI messages are shown in this table.</p> <table border="1"> <thead> <tr> <th>Exit Flag</th> <th>NBI Message</th> </tr> </thead> <tbody> <tr> <td>6</td> <td> <p>Some shadow solutions do not differ. Remove one of the non-competing objectives.</p> <p>Note CAGE does not run the subproblems if any pairs of shadow solutions are the same (within tolerance). All subproblems will show an exit flag of -8.</p> </td> </tr> <tr> <td>0</td> <td> <p>The solver stopped prematurely in at least one shadow problem and some shadow solutions do not differ.</p> <p>Note All shadow problems share the exit flag of 0, and CAGE does not run the subproblems (all will show an exit flag of -8)</p> </td> </tr> <tr> <td>-7</td> <td> <p>Solution is dominated by other solutions.</p> <p>Note Only successful solutions (with an <code>fmincon</code> exit flag ≥ 0) are used to determine whether a point is dominated.</p> </td> </tr> <tr> <td>-8</td> <td>NBI subproblem was not run because some shadow problems do not differ.</td> </tr> <tr> <td>-9</td> <td>NBI subproblem was not run because a shadow problem failed.</td> </tr> </tbody> </table>	Exit Flag	NBI Message	6	<p>Some shadow solutions do not differ. Remove one of the non-competing objectives.</p> <p>Note CAGE does not run the subproblems if any pairs of shadow solutions are the same (within tolerance). All subproblems will show an exit flag of -8.</p>	0	<p>The solver stopped prematurely in at least one shadow problem and some shadow solutions do not differ.</p> <p>Note All shadow problems share the exit flag of 0, and CAGE does not run the subproblems (all will show an exit flag of -8)</p>	-7	<p>Solution is dominated by other solutions.</p> <p>Note Only successful solutions (with an <code>fmincon</code> exit flag ≥ 0) are used to determine whether a point is dominated.</p>	-8	NBI subproblem was not run because some shadow problems do not differ.	-9	NBI subproblem was not run because a shadow problem failed.
Exit Flag	NBI Message												
6	<p>Some shadow solutions do not differ. Remove one of the non-competing objectives.</p> <p>Note CAGE does not run the subproblems if any pairs of shadow solutions are the same (within tolerance). All subproblems will show an exit flag of -8.</p>												
0	<p>The solver stopped prematurely in at least one shadow problem and some shadow solutions do not differ.</p> <p>Note All shadow problems share the exit flag of 0, and CAGE does not run the subproblems (all will show an exit flag of -8)</p>												
-7	<p>Solution is dominated by other solutions.</p> <p>Note Only successful solutions (with an <code>fmincon</code> exit flag ≥ 0) are used to determine whether a point is dominated.</p>												
-8	NBI subproblem was not run because some shadow problems do not differ.												
-9	NBI subproblem was not run because a shadow problem failed.												

See Also

Related Examples

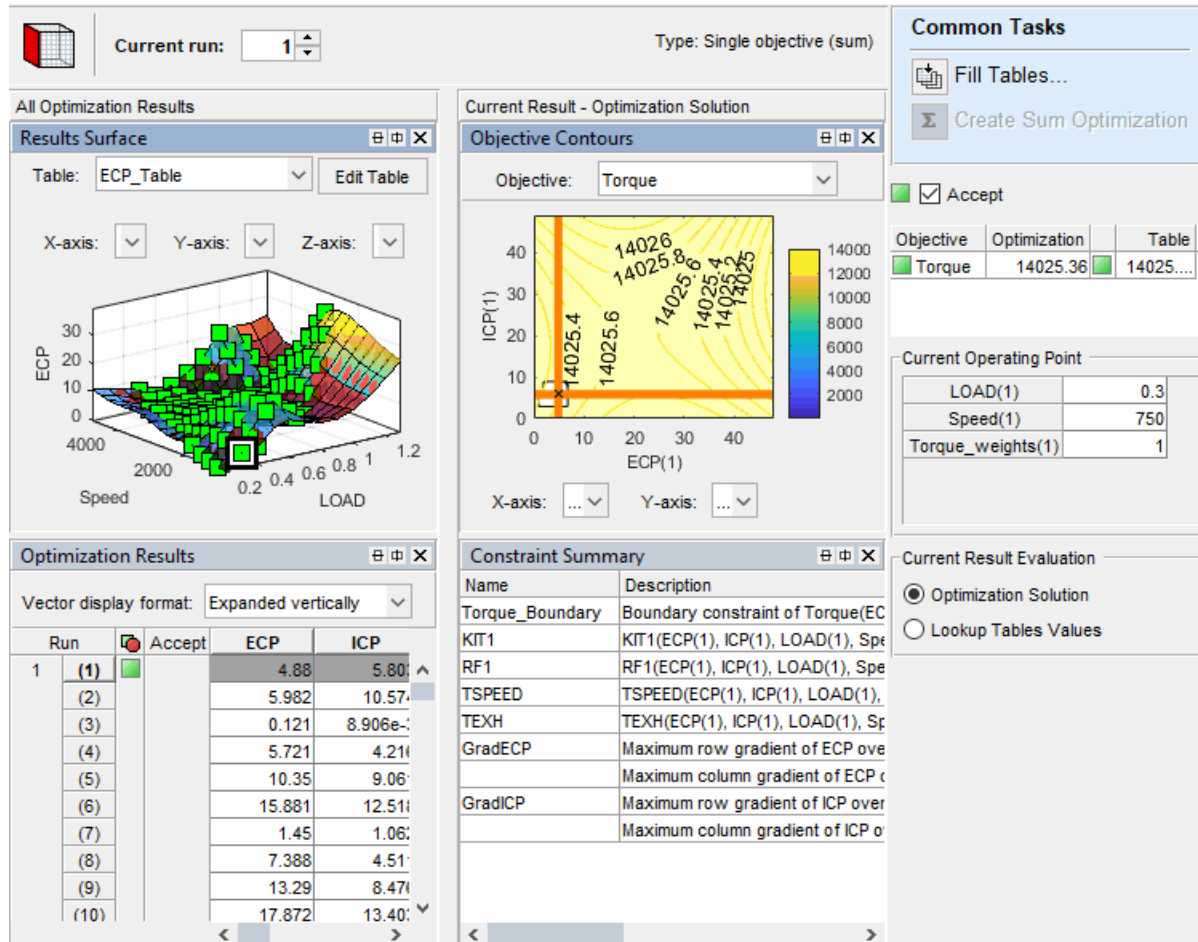
- “Examine the Multiobjective Optimization”

More About

- “gamultiobj Optimization Parameters” on page 7-49
- “NBI Optimization Parameters” on page 7-45
- “Selected Solution Slice” on page 8-39
- “Tools for Optimizations with Multiple Solutions” on page 8-38

Interpreting Sum Optimization Output

When you review sum optimization results, the default layout displays all the optimization results and results for specific solutions.




The left side of the layout shows the Results Surface and Optimization Results views for all the results. If you select a run, the:

- Objective Contours and Constraint Summary views provides the optimization results.
- Current Operating Pane shows the operating point.

In any view, you can change the view by right-clicking and selecting **Current View** to see the view options.

Objective Contour Plot

The objective contour plot for sum objective problems shows the contours of the objective. You can display plots against any pair of control parameters at each point in the set of operating points within each run. In this figure, a contour plot shows the results for run 129.



Constraint Summary

The constraint summary for sum optimizations shows a summary of all the constraint outputs for each constraint at the optimized control parameter settings for the selected run.

Name	Description	Constraint Value	Left Value	Right Value
Torque_Boundary	Boundary constraint of Torque(ECP(129), ICP(129), LOAD(129), Spee...	-2.852e-3	-2.852e-3	0
KIT1	KIT1(ECP(129), ICP(129), LOAD(129), Speed(129)) <= 2	-1.058	0.942	2
RF1	RF1(ECP(129), ICP(129), LOAD(129), Speed(129)) <= 25	-16.266	8.734	25
TSPEED	TSPEED(ECP(129), ICP(129), LOAD(129), Speed(129)) <= 200000	-1.594e5	40555.601	2e5
TEXH	TEXH(ECP(129), ICP(129), LOAD(129), Speed(129)) <= 860	-22.002	837.998	860
GradECP	Maximum row gradient of ECP over (LOAD,Speed)	-1.078	77.206	78.284
	Maximum column gradient of ECP over (LOAD,Speed)	-2.348e-4	0.021	0.021
GradICP	Maximum row gradient of ICP over (LOAD,Speed)	-16.941	61.342	78.284
	Maximum column gradient of ICP over (LOAD,Speed)	-5.386e-3	0.016	0.021

The rows of the table show a summary at each operating point. In this case, each of the rows corresponds to an evaluation of the constraint at each operating point within the run. For example,

the GradECP constraint rows provide an evaluation the maximum row and column gradients of ECP over load.

Optimization Results Table

Features of the Optimization Results table are labeled in this figure.

The screenshot shows a software interface titled "Optimization Output Values" with a "Vector display format" dropdown set to "Expanded vertically". The table below is a representation of the data shown in the screenshot, with red boxes highlighting the first row and the first five rows of the main data area. Arrows labeled A through G point to specific columns: A points to the Run index column, B to the quantity index column, C to the S, EXH, and INT columns, D to the Objective... column, E to the Objective1 column, F to the Constraint1-4 columns, and G to the table gradient constraint outputs.


Run	Accept	S	EXH	INT	Objective...	N	L	Objective1	Constraint1	Constraint2	Constraint3	Constraint4
1	<input checked="" type="checkbox"/>	5	5	5	5	5	5	1	5	5	24	24
(1)		15.122	22.347	39.162	1	1000	0.3	115.005	-153.321	-2.966e-8	-6.635e-3	-6.302e-3
(2)		17.923	15.784	44.132	1	1100	0.2		-184.163	-3.194e-8	-6.741e-3	-6.556e-3
(3)		18.565	24.801	41.214	1	1250	0.31		-151.336	-4.808e-11	-7.015e-3	-6.345e-3
(4)		22.818	19.487	44.162	1	1500	0.25		-162.578	-0.955	-7.124e-3	-6.513e-3
(5)		27.164	17.21	46.515	1	1625	0.18		-165.434	9.012e-9	-6.923e-3	-7.331e-3
(6)											-7.088e-3	-6.311e-3
(7)											-0.015	-0.016
(8)											-0.015	-0.015
(9)											-0.015	-0.016
(10)											-0.015	-0.015

Key to Optimization Results Table

- A: The run index — Index into the set of operating points that is being displayed.
- B: The quantity index.
 - For fixed and free variables, corresponds to the index of the operating point within the run.
 - For objectives, corresponds to the index of the output for the specific labeled objective.
 - For constraints, corresponds to the index of the output for the specific labeled constraint.
- C: Optimal Free Variable Settings — The optimal settings in this case of S, EXH, and INT at each operating point in the run. For example, the optimal settings of S, EXH and INT at the third operating point in this run 1 are S=18.565°, EXH= 24.801°, INT= 41.214°
- D: Fixed Variable Settings — Settings define the operating points for the run and other fixed variables (such as weights) required for objectives and constraints. These values were set up before the optimization was run.
- E: Optimal objective outputs — Optimal values of any objective outputs, that is, the optimized value of the weighted sum of TQ (115.002 Nm) over the five operating points shown in this case.
- F,G: Constraint outputs at optimized control parameter settings — The value of constraint outputs are displayed here. For the example problem, the model constraint outputs are displayed in the section labeled F. The number of constraint outputs matches the number of operating points. The table gradient constraint outputs are displayed in the section labeled G. The number of values returned by the table gradient constraint depends on the internal settings of that constraint.

Operating Point Indices


As in the Input Variable Values pane in the Optimization view, in the output view, the number in the bracket indicates the index of the operating point within a run.



In the Optimization Results table, the index of the operating point within the run is shown in brackets. In the Free Variable Values table and graphical displays, *InputVariableName(i)* provides the input variable at the *i*-th operating point within a run. For example, ECP(5) is the exhaust cam phase at the 5th operating point, and ICP(5) is the intake cam phase at the first operating point.


Objective Graphs

The objective graphs for sum objective problems show the objective cross section plots as in the point case. However, plots are now displayed against each control parameter at each point in the set of operating points within each run. In this figure, the weighted sum of TQ is plotted against the spark values at the first four operating points in run 1.



Constraint Graphs

The constraint graphs for sum objective problems show the cross section plots of the left side of the constraints as in the point case. However, in the sum case there are several more inputs and outputs that can be plotted. Specifically, each constraint can return several outputs. You can display the outputs against each control parameter at each point in the set of operating points within each run.




See Also

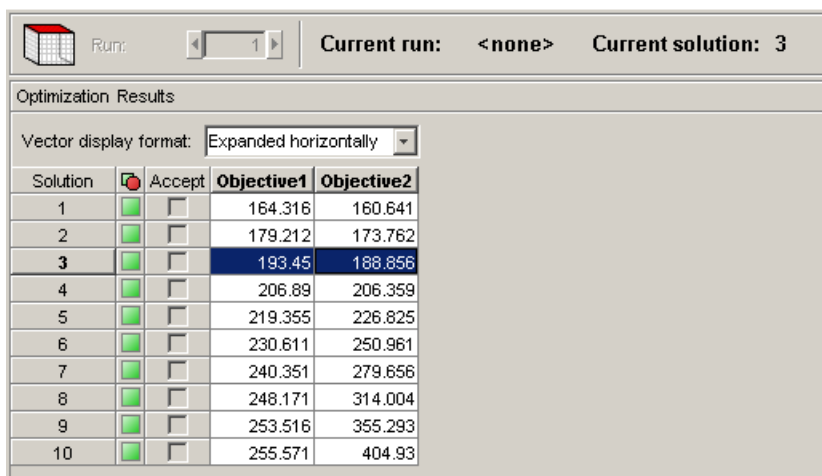
More About

- “Algorithm Restrictions” on page 7-18
- “Selecting Scalar Variables” on page 7-18
- “Solution Slice: Results Surface and Results Contour Views” on page 8-14

Weighted Objective Pareto Slice

The **Weighted Objective Pareto Slice** view (click ) shows a weighted sum Pareto solution. This table shows a weighted sum of the objective values over all runs for each solution. For a single objective optimization there is a single cell, which is the sum of the objective across all runs.

In this multiobjective example, the value in the **Objective1** column in the first row shows the sum of the solution 1 values of the first objective across all runs. The second row shows the sum of solution 2 **Objective1** values across all runs, and so on for all ten solutions in this case. This information can be useful, for example, for evaluating total emissions across a drive cycle. The default weights are unity (1) for each run.




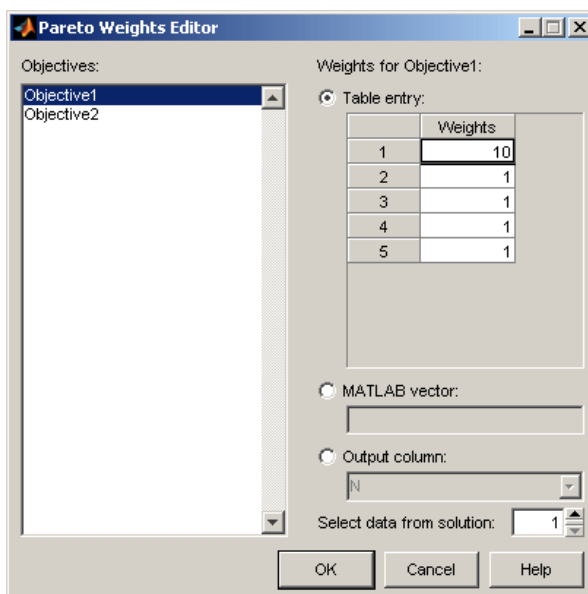
Run: Current run: <none> Current solution: 3

Optimization Results

Vector display format: Expanded horizontally

Solution	Accept	Objective1	Objective2
1	<input checked="" type="checkbox"/>	164.316	160.641
2	<input checked="" type="checkbox"/>	179.212	173.762
3	<input checked="" type="checkbox"/>	193.45	188.856
4	<input checked="" type="checkbox"/>	206.89	206.359
5	<input checked="" type="checkbox"/>	219.355	226.825
6	<input checked="" type="checkbox"/>	230.611	250.961
7	<input checked="" type="checkbox"/>	240.351	279.656
8	<input checked="" type="checkbox"/>	248.171	314.004
9	<input checked="" type="checkbox"/>	253.516	355.293
10	<input checked="" type="checkbox"/>	255.571	404.93

You can change the weights; for example, if you need a weighted sum of emissions over a drive cycle, you might want to give a higher weight to the value at idle speed. You can alter weights by clicking Edit Pareto Weights () in the toolbar. The **Pareto Weights Editor** appears.



Pareto Weights Editor

Objectives:

- Objective1
- Objective2

Weights for Objective1:

Table entry:

	Weights
1	10
2	1
3	1
4	1
5	1

MATLAB vector:

Output column:

Select data from solution:

OK Cancel Help

In this dialog box, you can select objectives to sum, and select weights for any run by clicking and editing, as shown in the previous example. The same weights are applied to each solution to calculate the weighted sums. Click **OK** to apply new weights, and the weighted sums are recalculated.

You can also specify weights with a MATLAB vector or any column in the optimization output by selecting the other option buttons. If you select **Output column** you can also specify which solution; for example, you could choose to use the values of spark from solution 5 at each operating point as weights. Click **Table Entry** again, and you can then view and edit these new values.

Note Weights applied in the **Weighted Pareto View** do not alter the results of your optimization as seen in other views. You can use the weighted sums to investigate your results only. You need to perform a sum optimization if you want to optimize using weighted operating points.

The Accept check box is disabled in this view. The exit flag is the minimum of all of the runs that are summed over, so the Accept status can only go green if all runs are green.

Writing User-Defined Optimizations

This section includes the following topics:

- “User-Defined Optimizations” on page 9-2
- “Create a User-Defined Optimization” on page 9-7
- “Creating an Optimization from Your Own Algorithm” on page 9-12
- “Optimization Function Reference” on page 9-23
- “Functions” on page 9-26

User-Defined Optimizations

In this section...

- “Introducing User-Defined Optimization” on page 9-2
- “Implementing Your Optimization Algorithm in CAGE” on page 9-2
- “About the Worked Example Optimization Algorithm” on page 9-4
- “Checking User-Defined Optimizations into CAGE” on page 9-6

Introducing User-Defined Optimization

User-defined optimizations are described in the following sections:

- “Implementing Your Optimization Algorithm in CAGE” on page 9-2 describes how to customize the optimization template to use your optimization routines in CAGE.
- There is a step-by-step guide to using the example provided to help you understand how to modify the template file to use your own optimization functions. See the tutorial section “Create a User-Defined Optimization” on page 9-7.

In many cases the standard routines supplied for constrained single objective (`fmincon`, `ga`, and `patternsearch`) and multiobjective optimization (`NBI`) are sufficient to allow you to solve your optimization problem. Sometimes, however, you need to write a customized optimization algorithm. This can be useful in many situations, for example,

- For an expert to capture an optimization process to solve a particular problem, for example, determination of optimal spark angle and exhaust gas recirculation rate on a port-fuel injection engine
- To implement an alternative optimization algorithm to those supplied
- To implement a complex constraint or objective that is only possible through writing code
- To produce custom output graphics

User-defined optimization functions in CAGE allow advanced users to write their own optimization routines that can access current CAGE data. In order to access the user function from CAGE, you must register the file with CAGE and place it on the MATLAB path. It is crucial that this function conforms to the template specified. The following sections describe this process.

Implementing Your Optimization Algorithm in CAGE

At some point a CAGE optimization function calls on an algorithm to optimize the objective functions over the free variables. You can implement the algorithm in the CAGE optimization function as an external MATLAB file. Use the template file as a basis for your optimization function. The best way to understand how to alter the template file to implement your own optimization algorithms is to compare it with the worked example, as described in the tutorial.

- See the following optimization tutorial sections:
 - “Create a User-Defined Optimization” on page 9-7 describes the process of using the worked example
 - “Creating an Optimization from Your Own Algorithm” on page 9-12 describes in detail the steps necessary to use an example optimization algorithm in CAGE

- “About the Worked Example Optimization Algorithm” on page 9-4, later on this page, examines the coding involved in implementing an external optimizer in a CAGE optimization file
- “Checking User-Defined Optimizations into CAGE” on page 9-6, later on this page, explains how to check in your optimization function so you can use it in CAGE

Optimization Function Structure

The optimization function files have two sections. To compare these sections in the worked example with the template file on which it is based:

- 1 Locate and open the file `mbc05template` in the `mbctraining` folder.
- 2 Type the following at the command line to open the example:

```
edit mbc05workedexample
```

The two sections are the `Options` section and `Evaluate` section.

- 1 The `Options` function section contains the settings that define your optimization. Here you can set up these attributes:
 - Name
 - Description
 - Free variables
 - Objective functions
 - Constraints
 - Helper data sets
 - Optimization parameters

CAGE interacts with the `cgoptimoptions` object, where all these settings are stored.

See “Methods of `cgoptimoptions`” on page 9-23 for information about setting up the options section.

If you leave the `cgoptimoptions` function unchanged, your optimization function must be able to support the default options. That is, your optimization will have:

- One objective
 - Any number of constraints (selected by the user in CAGE)
- 2 The `Evaluate` function section contains your optimization routine. CAGE calls this section when the **Run** button is clicked.

Place your optimization routine under this section, interacting with CAGE (obtaining inputs and sending outputs) via the `cgoptimstore` object. Your optimization must conform to the following syntax:

```
optimstore = <Your_Optimization> (optimstore)
```

where `<Your_Optimization>` is the name of your optimization function.

Any local functions called by your optimization routine should also be placed at the bottom of this section.

See “Methods of `cgoptimstore`” on page 9-24.

Note Be careful not to overwrite the worked example and template files when you are trying them out — save them under a new name when you make changes.

There is a step-by-step guide describing how to modify the template using the worked example optimization function in the optimization tutorial. See “Create a User-Defined Optimization” on page 9-7.

About the Worked Example Optimization Algorithm

`mbcweoptimizer` is an example of a user-specified optimization that solves the following problem:

max TQ over (AFR, SPK).

- `[bestafr, bestspk] = mbcweoptimizer(TQ)` finds a maximum (`bestafr, bestspk`) to the function TQ.

TQ must be a function (or a function handle) where the first two input arguments are AFR and SPK respectively. TQ functions with more parameters can be used. The extra parameters to these functions can be specified using anonymous functions. For example if a TQ model has N and L inputs, you can use the following call to `mbcweoptimizer`:

```
[bestafr, bestspk] = mbcweoptimizer(@(afr, spk)TQ(afr, spk, N, L))
```

- `[bestafr, bestspk]=mbcweoptimizer(TQ, afrrng, spkrng)` finds a maximum (`bestafr,bestspk`) to the function TQ.

`afrrng` and `spkrng` are 1-by-2 row vectors containing search ranges for those variables.

- `[bestafr, bestspk]=mbcweoptimizer(TQ, afrrng, spkrng, res)` finds a maximum (`bestafr,bestspk`) to the function TQ.

This optimization is performed over a `res`-by-`res` grid of (AFR, SPK) values. If `res` is not specified, the default grid resolution is 25.

The Structure of the Worked Example

The best way to understand how to implement an external optimizer in a CAGE optimization function is to study the details of the example.

- To view the whole worked example file, at the command line, type

```
edit mbc05workedexample
```

The following code section is taken from the Evaluate section of the worked example file as an example.


```

78 % For every fixed point, find the optimum (afr, spk) using
79 % the mbcweoptimizer routine you have written
80 - [bestafr, bestspk] = mbcweoptimizer,@n_evalTQ, [minAFR, maxAFR], ...
81     [minSPK, maxSPK], res);
82
83 % Set the best values calculated for the free variable(s) into the
84 % data set
85 - optimstore = setFreeVariables(optimstore, [bestafr, bestspk]);
86
87 % Return some information about the optimization
88 - OUTPUT.Algorithm = 'Brute force search';
89 - OUTPUT.Resolution = res;
90
91 % Set all information in the optimstore
92 - optimstore = setExitStatus(optimstore, 1, 'Optimization Completed');
93 - optimstore = setOutput(optimstore, OUTPUT);

```

The code fragment above is in the `i_Evaluate` local function. This local function is called once for each run of the script. The line of code labeled A above calls the worked example optimization algorithm external to the optimization function. As with functions in the Optimization Toolbox product, the first argument to the call to the optimizer is a function handle that evaluates the objectives at a given input point. We recommend you place the function pointed at by the function handle in the optimization file. If you do not place them in the same file you must make sure the `evaluate` function file is on the MATLAB path. As an example, the optimization evaluation function in the worked example optimization is shown in the code fragment following.

```

*****
% Objective evaluation function
*****
function y = n_evalTQ(afr, spk)

    y = evaluate(optimstore, [afr, spk]);

end

```

The inputs to `n_evalTQ` are the required inputs for the torque (in this case) model. To evaluate the objective, the `evaluate` method from the `optimstore` object is used. In the above example, the line of code referenced by B evaluates the torque model in the worked example at the `(afr, spk)` input points. The values of `(N, L)` at the current run are used in the evaluation of the torque model. CAGE retrieves these values from `optimstore` when the torque model is evaluated.

The two local functions presented above are an example of how to implement an external optimizer in a CAGE optimization file.


See also the optimization tutorial section “Creating an Optimization from Your Own Algorithm” on page 9-12, which describes in detail the steps involved in incorporating an example algorithm into a CAGE optimization file.

Checking User-Defined Optimizations into CAGE

When you have modified the template to create your own optimization function, you must check it into the Model-Based Calibration Toolbox product in order to use the function in CAGE. Once you have checked in your optimization function it appears in the **Optimization Wizard**. See “Optimization Wizard” on page 7-10.

To check a user-defined optimization into CAGE,

- 1 Select **File -> Preferences**.
- 2 Click the **Optimization** tab and click **Add...** to browse to your file. Select the file and click **Open**. This registers the optimization function with CAGE. You need to do this when you customize your own optimizations.



The example shows the worked example function, which is already registered with CAGE for use in the optimization tutorial.

- 3 You can click **Test** to check that the optimization function is correctly set up. This is a very useful function when you use your own functions; if anything is incorrectly set up the test results tell you where to start correcting your function.

You can see an example of this by saving a copy of the worked example file and changing one of the variable names (such as `afr`) to a number. Try to check this altered function into CAGE and the **Test** button will return an informative error specifying the line you have altered.

- 4 Click **OK** to dismiss the **CAGE Preferences** dialog box and return to the CAGE browser.

Registered optimizations appear in the **Optimization Wizard** when you set up a new optimization.

Registered optimizations appear in the **Create Optimization from Model Wizard** *unless* your user-defined optimization script defines operating point sets and/or a fixed number of free variables. This is common with Version 2.0 scripts. If this is the case you must use the **Optimization Wizard** instead.

Create a User-Defined Optimization

In this section...

“Overview” on page 9-7

“Using the Worked Example Optimization” on page 9-8

Overview

There is a simple worked example provided to show you what you can do by modifying the template file to write your own optimizations. This example demonstrates a simple use of the CAGE optimization feature. The aim of this example is to obtain values of spark (SPK) and air/fuel ratio (AFR) that maximize torque at a given speed (N) and load (L). These values could then be used to fill calibration tables.

An example of a user-defined optimization algorithm is provided.

- To see a description of this algorithm, at the command line type

```
help mbcweoptimizer
```

`mbcweoptimizer` is an example of a user-specified optimization that solves the following problem:

Maximum TQ over (AFR, SPK) at a given (N, L) point.

The syntax for this example function, `mbcweoptimizer`, mimics that used in the Optimization Toolbox product.

- To evaluate this at the command line, type this example:

```
[bestafr, bestspk] = mbcweoptimizer(@(afr, spk)mbcTQ(afr,...
spk, 1000, 0.2))
```

The optimization finds values of AFR and spark (the free variables) that give the maximum output from TQ at the values of speed and load (the fixed variables) that you specified, in this case speed = 1000, load = 0.2, as shown below.

```
bestafr =
12.9167
bestspk =
25
```

To use this optimization algorithm in CAGE, you need to include the function in a CAGE optimization function script. This worked example modifies the template provided to show you how to use your own algorithms within CAGE.

- To view the worked example file, at the command line, type

```
edit mbc0Sworkedexample
```

The worked example optimization wraps `mbcweoptimizer` in a function that can be called by the CAGE optimization feature. When you run your optimization from CAGE, you can alter the search ranges of the free variables and the resolution of the search.

Using the Worked Example Optimization

In order to run any optimization, you first need to set up your CAGE session with a model.

For this example, the CAGE session requires


- A torque model
- A variable dictionary defining required variable ranges and set points (N, L, AFR, and SPK)
- A data set defining the (N,L) operating points where you want to run the optimizer

There is a preconfigured session provided that contains the model, variable dictionary, and data set.

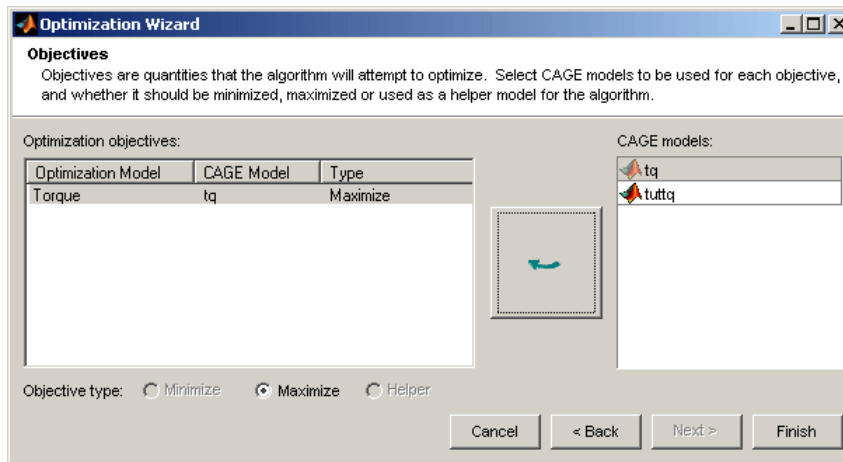
- 1 Select **File > Open Project** and load the file `optimworkedexample.cag`. This is in the `mbctraining` folder.
 - You can look at the variables by clicking the **Variable Dictionary** button in the **Data Objects** pane.
 - You can look at the operating point set by clicking **Data Sets** in the **Data Objects** pane. Note you can specify fixed variables for optimizations either directly in the optimization view or import them from a data set or table.
- 2 Select **File > New > Custom Optimization**.

The Optimization Wizard appears.

- 3 Select `WorkedExample`, and click **Next**.
- 4 Associate each pair of inputs and variables, by clicking `afr` and `A` in the left and right lists, and then click the **Select** button. Similarly associate `spark` with `spk`. Click **Next**.

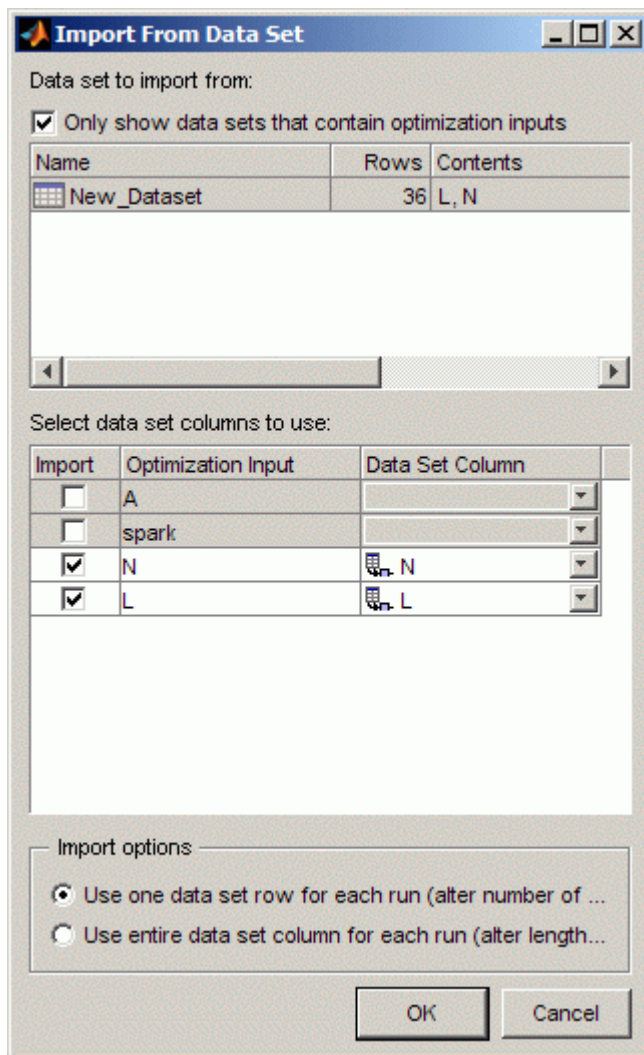


- 5 The next screen of the wizard automatically shows the **Torque** model selected and **Maximize** chosen; these are specified in the function. Select `tq` in the CAGE model list and click the button to match it with the Torque optimization model, then click **Finish**.



CAGE switches to the **Optimization** view and the new **Optimization** node appears in the tree.

- 6 If you ran the optimization now it would run at one point, the set point of all the variables. You use the free and fixed **Variable Values** panes to select operating points. You can edit points manually or import them. Select **Optimization > Import From Data Set**.

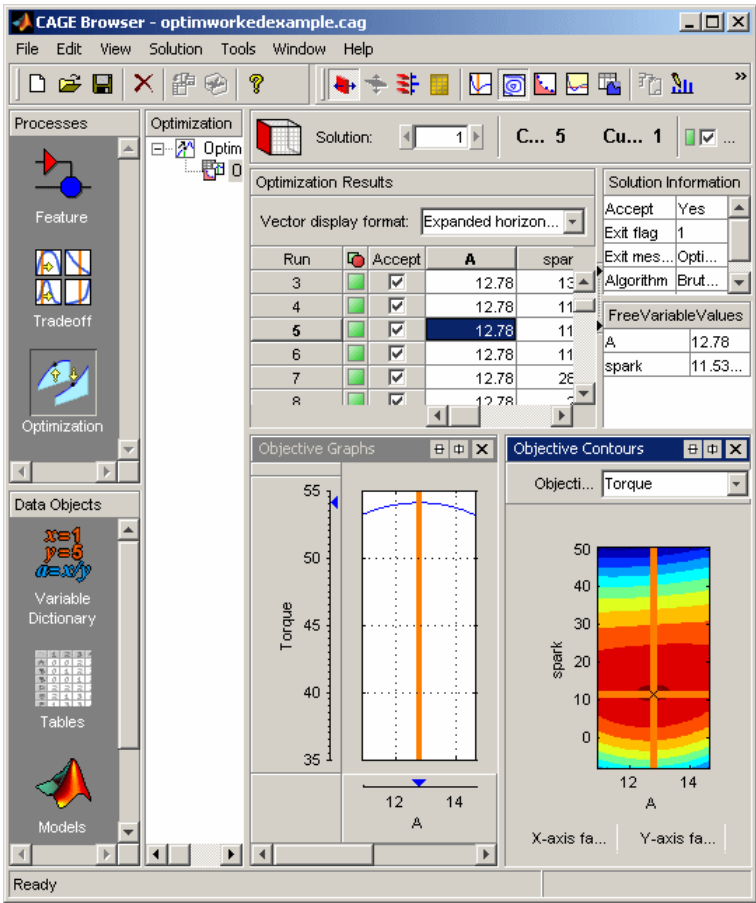


The project file contains a data set with N and L values, and these are automatically selected. Click **OK** to import.

Notice 36 rows appear in both fixed and free variable panes, and operating point values have been imported into the N and L columns in the **Fixed Variables** pane. The initial values for A and spark for each point are the set points in the variable dictionary.

- 7 Click Run Optimization in the toolbar.
- 8 When the optimization completes, the view switches to the new `Optimization_Output` node.

The output display should look like the following. The optimization has found the values of SPK and AFR that give the maximum model value of torque at each operating point specified. Select different operating points by clicking in the table: the model plots at the selected operating point are shown. There is only one solution per operating point, so you cannot scroll through the solutions.



See Also

More About

- “Creating an Optimization from Your Own Algorithm” on page 9-12

Creating an Optimization from Your Own Algorithm

In this section...

“Process Overview” on page 9-12
 “Step 1: Verify the Algorithm” on page 9-13
 “Step 2: Create a CAGE Optimization Function” on page 9-14
 “Step 3: Define the Optimization Options” on page 9-15
 “Step 4: Add the Algorithm to the Optimization Function” on page 9-17
 “Step 5: Register Your Optimization Function with CAGE” on page 9-19
 “Step 6: Verify Your New Optimization” on page 9-20

Process Overview

The CAGE optimization feature allows you to use your own optimization algorithms as alternatives to the library routines `fmincon`, `NBI`, `ga` and `patternsearch`.

Using an example, this tutorial illustrates how to take an existing optimization algorithm and implement it as an optimization function for use in CAGE optimization.

The problem to be solved is the worked example problem:

Maximize torque (TQ) over the free variables (SPK, AFR) over a specified set of (N, L) points. These points are defined in the data set `New_Dataset`, which can be found in the CAGE session `optimworkedexample.cag` and can be imported to the fixed variable values pane in the Optimization view.

The torque model to be used is that in `/mbctraining/Holiday.mat`.

The process steps are:

- 1 Start with your own algorithm. We will use `fminunc` from the Optimization Toolbox product as an example.
- 2 Create a CAGE optimization function.
- 3 Define the attributes of your optimization in the CAGE optimization function.
- 4 Add your algorithm to the CAGE optimization function.
- 5 Register your completed optimization function with CAGE.
- 6 Verify the optimization.

The steps of this tutorial lead you through a series of examples illustrating how to construct the code to incorporate your own algorithm into an optimization in CAGE.

Before you begin you must create a working folder.

- 1 Create a new folder (for example, `C:\Optimization_Work`). We recommend that you place this folder outside your MATLAB folders to avoid interfering with toolbox files.
- 2 Copy the following six files from the `mbctraining` folder into your new working folder:

```
curr_tutoptim.m
mbc0Stemplate.m
```



```

mbc0Stutoptimfunc_s1.m
mbc0Stutoptimfunc.m
optimtut.mat
optimtuteg.mat

```

- 3 Make sure your new working folder is on the MATLAB path; either change **Current Folder** in MATLAB to the new working folder, or add the folder to the path as follows:
 - a On the **Home** tab, in the **Environment** section, click **Set Path**.
 - b Click **Add Folder** and browse to your working folder.
 - c Click **OK**.
 - d Click **Save**.
 - e Click **Close**.

Step 1: Verify the Algorithm

`curr_tutoptim.m` is an example file to verify that `fminunc` solves the worked example problem. You can try this at the MATLAB command line.

- 1 To open the algorithm file in the Editor, either enter `open curr_tutoptim.m` at the command line, or if the **Current Folder** in MATLAB is your new working folder, then double-click `curr_tutoptim.m` in the **Current Folder**. You should see the code in the MATLAB editor.
- 2 To verify that `fminunc` solves the worked example problem, type the following command at the MATLAB prompt:

```
bestX = curr_tutoptim
```

After the progress messages complete the workspace output should resemble the following:

```

BestX =
    23.768    12.78
    18.179    12.78
    14.261    12.78
    12.014    12.78
    11.439    12.78
    12.535    12.78
    27.477    12.78
    21.887    12.78
    17.969    12.78
    15.722    12.78
    15.147    12.78
    16.243    12.78
    31.185    12.78
    25.595    12.78
    21.677    12.78
     19.43    12.78
    18.855    12.78
    19.951    12.78
    34.893    12.78
    29.303    12.78
    25.385    12.78
    23.138    12.78
    22.563    12.78
    23.659    12.78
    38.601    12.78

```

33.012	12.78
29.093	12.78
26.847	12.78
26.271	12.78
27.368	12.78
42.309	12.78
36.72	12.78
32.802	12.78
30.555	12.78
29.979	12.78
31.075	12.78

The matrix `bestX` contains the optimal SPK and AFR values that maximize the MBC model torque (exported from `Holliday.mat`) at the speed and load points defined in the matrix data.

`fminunc` is the example optimization algorithm that you want to transfer to CAGE for use in the optimization GUI.

This tutorial shows how to make `fminunc` available for use in the CAGE optimization feature.

Step 2: Create a CAGE Optimization Function

Any optimization algorithm you want to use in CAGE must be contained in an optimization function. A CAGE optimization function consists of two sections.

The first section defines the following attributes of the optimization:

- A name for the optimization
- A description of the optimization
- Number of free variables
- Labels for free variables (if required), so the user can match variables in CAGE to the required algorithm free variables.
- Number of objectives
- Labels for objective functions, so the user can match models in CAGE to the required algorithm objectives (you can match in CAGE, so labels do not have to be exact in the optimization function)
- Number of constraints
- Labels for constraints, so the user can match models in CAGE to the required models in your algorithm constraints
- Number of data sets
- Labels for data sets, so the user can match data sets in CAGE to the required variable data for your algorithm
- Any other parameters required by the optimization algorithm

The second section contains the optimization algorithm.

Open `mbc0Stemplate.m` in the MATLAB editor.

`mbc0Stemplate.m` is an empty CAGE optimization function. The two (currently empty) sections of the function are `options` (for defining optimization attributes) and `optimstore` (for defining your optimization algorithm). Note that this file can be used as a template for any optimization function that you write.

Step 3: Define the Optimization Options

The next step is to define the attributes of your optimization (in Section 1 of the template).

Open `mbc0Stutoptimfunc_s1.m`. In this file, you can see the optimization attributes that have been defined.

The following is a code fragment from this file:

```
% Deal with the action inputs
if strcmp(action, 'options')

    options = in;
    % Add a name
    options = setName(options, 'Tutorial_Optimization');

    % Add a description
    options = setDescription(options, 'A simple worked example to maximize torque')

    % Set up the free variables
    options = setFreeVariablesMode(options, 'fixed');
    options = addFreeVariable(options, 'afr');
    options = addFreeVariable(options, 'spk');

    % Set up the objective functions
    options = setObjectivesMode(options, 'fixed');
    options = addObjective(options, 'Torque', 'max');

    % Set up the constraints
    options = setConstraintsMode(options, 'fixed');
    % There are no constraints for this example

    % Set up the operating point sets
    options = setOperatingPointsMode(options, 'fixed');
    % There are no operating point sets for this example

    % Set up the optimization parameters
    options = addParameter(options, 'Display', ...
        {'list', {'none', 'final', 'iter'}}, 'none');
    options = addParameter(options, 'MaxIter', ...
        {'integer', 'positive'}, 200, 'Maximum iterations');
    options = addParameter(options, 'MaxFunEvals', ...
        {'integer', 'positive'}, 1000, 'Maximum function evaluations');
    options = addParameter(options, 'TolX', ...
        {'number', 'positive'}, 1e-6, 'Variable tolerance');
    options = addParameter(options, 'TolFun', ...
        {'number', 'positive'}, 1e-6, 'Function tolerance');

    out= options;

elseif strcmp(action, 'evaluate')

    optimstore = in;

    %
    % Put optimization algorithm here
    %
```

The optimization attributes are passed to CAGE via the `cgoptimoptions` object, referenced by `options` in the code in `mbc0Stutoptimfunc_s1.m`. See after the table for details of the `cgoptimoptions` object. The `cgoptimoptions` object has a set of functions that set the optimization attributes in CAGE. This is where you specify the name, description, free variables, objective functions, constraints, helper data sets, and optimization parameters for the optimization.

For detailed information on all the available functions, see “Optimization Function Reference” on page 9-23 in the CAGE documentation. The above code has used the `cgoptoptions` object (`options`) to set the optimization attributes as described in the following table.

Look through the code to locate the listed **Code Section Where Set** for each attribute to see how each of the optimization options is set up.

Attribute	Value	Code Section Where Set
Optimization Name	Tutorial_Optimization	Add a name - setName
Description	A simple worked example to maximize torque	Add a description - setDescription
Number of Free Variables	Cannot be changed by the user in the GUI (the mode has been set to 'fixed')	Set up the free variables - setFreeVariablesMode
Required Free Variables	This function requires two free variables, labeled 'afr' and 'spk'. The user matches these free variable labels to CAGE variables in the Optimization Wizard.	Set up the free variables - addFreeVariables
Number of Objectives	Cannot be changed by the user in the GUI (the mode has been set to 'fixed')	Set up the objective functions - setObjectivesMode
Required Objective functions	This function requires one objective function, which will be labeled 'Torque' in the optimization feature. The user matches this 'Torque' label to a CAGE model.	Set up the objective functions - addObjective
Number of Constraints	Cannot be changed by the user in the GUI (the mode has been set to 'fixed')	Set up the constraints - SetConstraintsMode
Required Constraints	As the mode is fixed and no constraint labels have been defined, this optimization has no linear or nonlinear constraints.	Set up the constraints - %There are no constraints
Number of Helper Data Sets	Cannot be changed by the user in the GUI (the mode has been set to 'fixed'). There are no helper data sets for this example.	Set up the operating point sets - setOperatingPointsMode
Optimization Parameters	This function will allow the user to change five parameters. These will be displayed in the Optimization Parameters dialog box and labeled Display, Maximum iterations, Maximum function evaluations, Variable tolerance, and Function tolerance.	Set up the optimization parameters - addParameter

When one of your optimizations is created in the CAGE GUI, CAGE first calls your optimization function to define the attributes of the optimization. The function call from CAGE has the form

```
optionsobj = <your_optimization_function>('options', optionsobj)
```

This is how your optimization function receives the `cgoptoptions` object. Note that your optimization function *must* support this interface.

Step 4: Add the Algorithm to the Optimization Function

In this step you complete the optimization function by adding your algorithm. To do this, a few changes need to be made to the code that calls the algorithm, as data (for example, free variable values, constants, and so on) will now be passed to and from CAGE rather than from the MATLAB workspace.

1 Open `mbc0Stutoptimfunc.m`.

This file contains the completed optimization algorithm. The following is a code fragment from this file.

```
elseif strcmp(action, 'evaluate')
    optimstore = in;
    optimstore = tutoptimizer(optimstore);
    out = optimstore;
```

A single line has been added, namely

```
optimstore = tutoptimizer(optimstore)
```

This line calls the modified optimization algorithm. Note the syntax of the algorithm: it *must* take the form

```
optimstore = <your_optimization_algorithm>(optimstore)
```

2 The local function `tutoptimizer` can be found at the bottom of the file `mbc0Stutoptimfunc.m`. Scroll down to view the algorithm, modified for use in CAGE.

`optimstore` is a `cgoptimstore` object. This is an interface object that allows you to get data from and set data in the CAGE optimization feature. You can now see how the `optimstore` object is used by comparing the modified optimization algorithm, `tutoptimizer`, with the original algorithm, `curr_tutoptim`, for each of the main sections of the algorithm.

The following sections illustrate how to convert an existing algorithm for use in CAGE. Note that in this tutorial example, the code is already modified for you to examine.

Algorithm Section 1

Get the start conditions (x_0) for the free variables.

Original code:

x_0 passed in from the MATLAB workspace.

Modified code:

```
 $x_0$  = getInitFreeVal(optimstore);
```

In the original algorithm, x_0 is passed into the algorithm from the MATLAB workspace. In CAGE, we invoke the `getInitFreeVal` function on the `optimstore` object to retrieve x_0 .

Algorithm Section 2

Perform the optimization (in Section 2 of the template).

Original code (from `curr_tutoptim`):

```
[bestx(i, :), notused1, notused2, OUTPUT(i)] = fminunc(trqfunc,
x0, algoptions);
```

which calls the following code to evaluate the cost function:

```
function tq = trqfunc(x)

    % Evaluate torque. Note x = [SPK, AFR]
    tq = EvalModel(TQMOD, [x(1), N(i), L(i), x(2)]);

    % Maximising torque, so need to return -tq
    tq = -tq;

end
```

Modified code:

```
[bestx, unused, exitFlag, OUTPUT] = fminunc(@trqfunc_new,
x0, algoptions);
```

which calls the following code to evaluate the cost function:

```
function y = trqfunc_new(x)
    % Evaluate the torque objective function
    y = -evaluate(optimstore, x);
end
```

In performing the algorithm, the only difference between the original and modified code is how the objective function is evaluated. The original algorithm requires the objective function (a Model-Based Calibration Toolbox model for torque) to be loaded in and evaluated as required. In the modified algorithm the objective function (torque) is evaluated by invoking the `evaluate` function on the `optimstore` object. Note that the inputs to the torque model are passed in to the `evaluate` function as shown in the following table.

Original Input	Input to Evaluate Function
S	X(1)
A	X(2)

Algorithm Section 3

Retrieve output data.

Original code:

Optimal free variable settings are returned through the variable `bestX` in `curr_tutoptim`.

Modified code:

```
% Write results to the optimstore
optimstore = setFreeVariables(optimstore, bestx);

% Set termination message
termMsg = OUTPUT.message;
OUTPUT = rmfield(OUTPUT, 'message');

% Set all information in the optimstore and leave
```

```
optimstore = setExitStatus(optimstore, exitFlag, termMsg);
optimstore = setOutput(optimstore, OUTPUT);
```

In the modified algorithm, the results need to be sent back to the CAGE optimization feature and not the MATLAB workspace. To do this, optimization results are set in the `optimstore` object, which is then returned to CAGE. There are three functions you should invoke on the `optimstore` object to return optimization results to CAGE:

- `setFreeVariables` — Returns the optimal free variable values to CAGE
- `setExitStatus` — Returns an integer that indicates whether the algorithm terminated successfully or not (positive is successful). This sets the termination message.
- `setOutput` — Returns any diagnostic information on the algorithm to CAGE

Step 5: Register Your Optimization Function with CAGE

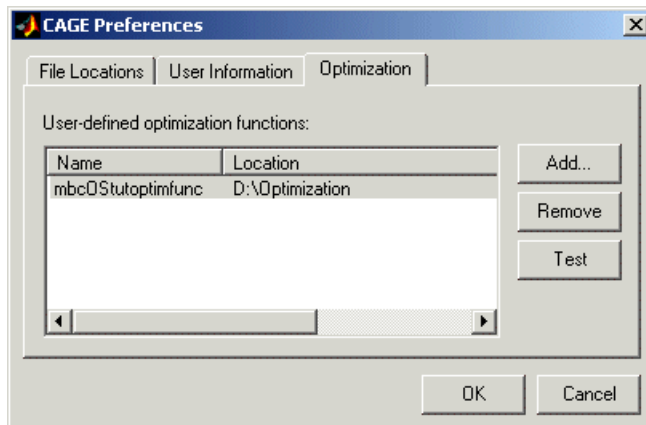
The worked example provided is preregistered so you can see it as an option in the Optimization Wizard when setting up a new optimization. You must register new functions before you can use them. When you have modified the template to create your own optimization function, as in this example, you must register it with the Model-Based Calibration Toolbox product in order to use the function in CAGE. Once you have checked in your optimization function it appears in the Optimization Wizard.

- 1 In CAGE, select **File > Preferences**.

The CAGE Preferences dialog appears.

- 2 Click the **Optimization** tab and click **Add** to browse to your file.
- 3 Locate the file `mbc0Stutoptimfunc.m` (in the working folder you created) and click **Open**.

This registers the optimization function with CAGE.



- 4 You can now test the function by clicking **Test**. This is a good check for any syntax errors in your optimization function. This is a very useful function when you use your own functions; if anything is incorrectly set up the test results will tell you where to start correcting your function.

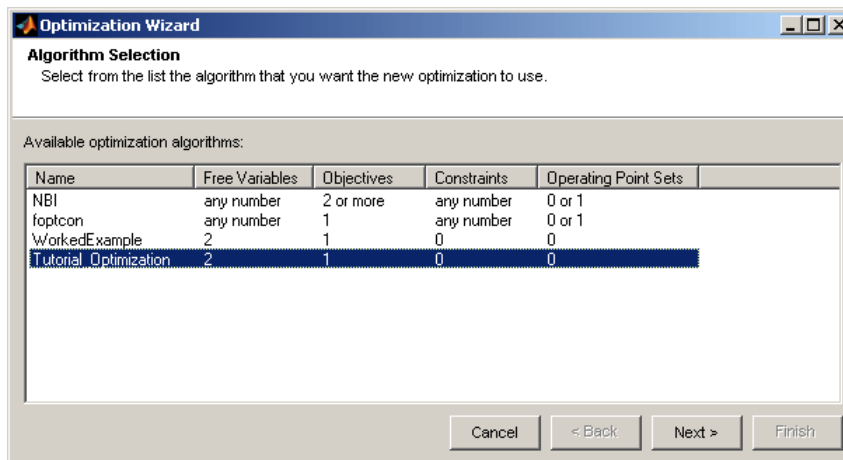
You could see an example of this by saving a copy of the worked example file and changing one of the variable names (such as `afr`) to a number. Try to check this altered function into CAGE, and the **Test** button will return an informative error specifying the line you have altered.

- Click **OK** to leave the CAGE Preferences dialog. If the optimization function tested successfully, it is registered as an optimization function that can be used in CAGE, and appears in the Optimization Wizard.


Step 6: Verify Your New Optimization

To verify the algorithm we set up a CAGE session to run the optimization that was performed in step 1. For this example, the CAGE session has already been set up. Follow the steps below to run the tutorial optimization in CAGE.

- In CAGE, select **File > Open Project** and load the file `optimworkedexample.cag` (unless you already have this project open). This project is in the `mbctraining` folder.
- Select **File > New > Optimization**.
- The newly registered optimization appears in the list of algorithm names. Select `Tutorial_Optimization` from the list. Click **Next**.

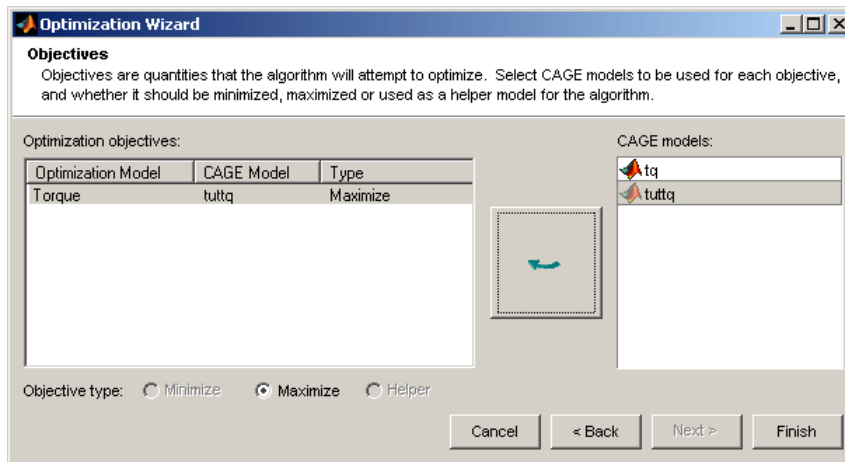


- Match the variables as shown.



Click **Next**.

- Match the Torque model to the `tuttq` CAGE model as shown.



Click **Finish**.






































- 6 If you ran the optimization now it would run at one point, the set point of all the variables. You use the free and fixed **Variable Values** panes to select operating points. You can edit points manually or import them. Do one of the following:
 - If you have the previous worked example optimization in your current session, in the optimization view increase the **Number of runs** to 36, and then copy and paste the fixed variable values from the previous optimization.
 - If you do not have the previous optimization in your session, select **Optimization > Import From Data Set**. The project file contains a data set with N and L values, and these are automatically selected. Click **OK** to import.

Now you should have 36 rows in both fixed and free variable panes, and operating point values in the N and L columns in the **Fixed Variables** pane. The initial values for A and spark for each point are the set points in the variable dictionary.

- 7 Select **Optimization > Set Up**. The Optimization Parameters dialog box appears. Observe the five parameters defined in the tutorial optimization script.

Change the variable and function tolerances to $1e-4$, and click **OK** to close the dialog box.

- 8 Run the optimization and view the results. The output data matrix should resemble the following. Note that the optimal values for A and SPK are very similar to those from the original algorithm.

Run	 Accept	A	spark	N	L	Torque
1	 <input checked="" type="checkbox"/>	12.78	23.768	1000	0.2	9.07
2	 <input checked="" type="checkbox"/>	12.78	18.179	1000	0.3	20.319
3	 <input checked="" type="checkbox"/>	12.78	14.261	1000	0.4	31.567
4	 <input checked="" type="checkbox"/>	12.78	12.014	1000	0.5	42.816
5	 <input checked="" type="checkbox"/>	12.78	11.441	1000	0.6	54.064
6	 <input checked="" type="checkbox"/>	12.78	12.534	1000	0.7	65.313
7	 <input checked="" type="checkbox"/>	12.78	27.476	2000	0.2	9.742
8	 <input checked="" type="checkbox"/>	12.78	21.887	2000	0.3	20.99
9	 <input checked="" type="checkbox"/>	12.78	17.969	2000	0.4	32.238
10	 <input checked="" type="checkbox"/>	12.78	15.722	2000	0.5	43.487
11	 <input checked="" type="checkbox"/>	12.78	15.147	2000	0.6	54.735
12	 <input checked="" type="checkbox"/>	12.78	16.243	2000	0.7	65.984
13	 <input checked="" type="checkbox"/>	12.78	31.185	3000	0.2	9.342
14	 <input checked="" type="checkbox"/>	12.78	25.597	3000	0.3	20.591
15	 <input checked="" type="checkbox"/>	12.78	21.683	3000	0.4	31.839
16	 <input checked="" type="checkbox"/>	12.78	19.431	3000	0.5	43.087
17	 <input checked="" type="checkbox"/>	12.78	18.856	3000	0.6	54.336
18	 <input checked="" type="checkbox"/>	12.78	19.953	3000	0.7	65.584
19	 <input checked="" type="checkbox"/>	12.78	34.891	4000	0.2	7.872
20	 <input checked="" type="checkbox"/>	12.78	29.305	4000	0.3	19.12
21	 <input checked="" type="checkbox"/>	12.78	25.385	4000	0.4	30.369
22	 <input checked="" type="checkbox"/>	12.78	23.139	4000	0.5	41.617
23	 <input checked="" type="checkbox"/>	12.78	22.563	4000	0.6	52.866
24	 <input checked="" type="checkbox"/>	12.78	23.659	4000	0.7	64.114
25	 <input checked="" type="checkbox"/>	12.78	38.614	5000	0.2	5.331
26	 <input checked="" type="checkbox"/>	12.78	33.013	5000	0.3	16.58
27	 <input checked="" type="checkbox"/>	12.78	29.093	5000	0.4	27.828
28	 <input checked="" type="checkbox"/>	12.78	26.847	5000	0.5	39.077
29	 <input checked="" type="checkbox"/>	12.78	26.27	5000	0.6	50.325
30	 <input checked="" type="checkbox"/>	12.78	27.367	5000	0.7	61.574
31	 <input checked="" type="checkbox"/>	12.78	42.321	6000	0.2	1.72
32	 <input checked="" type="checkbox"/>	12.78	36.721	6000	0.3	12.969
33	 <input checked="" type="checkbox"/>	12.78	32.801	6000	0.4	24.217
34	 <input checked="" type="checkbox"/>	12.78	30.555	6000	0.5	35.465
35	 <input checked="" type="checkbox"/>	12.78	29.98	6000	0.6	46.714
36	 <input checked="" type="checkbox"/>	12.78	31.075	6000	0.7	57.962

Optimization Function Reference

In this section...
“Methods of cgoptimoptions” on page 9-23
“Methods of cgoptimstore” on page 9-24

Methods of cgoptimoptions

You use these functions to set up all your optimization settings in the `Options` section of the file. You can set up any or all of these seven attributes:

- Name
- Description
- Free variables
- Objective functions
- Constraints
- Helper data sets
- Optimization parameters

The following methods are available:

<code>addFreeVariable</code>	Add free variable to optimization
<code>addLinearConstraint</code>	Add linear constraint to optimization
<code>addModelConstraint</code>	Add model constraint to optimization
<code>addObjective</code>	Add objective to optimization
<code>addOperatingPointSet</code>	Add operating point set to optimization
<code>addParameter</code>	Add parameter to optimization
<code>getConstraints</code>	Return information about all optimization constraints
<code>getConstraintsMode</code>	Return current usage of constraints
<code>getDescription</code>	Get current description for optimization function
<code>getEnabled</code>	Get current enabled status for optimization
<code>getFreeVariables</code>	Return optimization free variable labels
<code>getFreeVariablesMode</code>	Return current usage of free variables
<code>getLinearConstraints</code>	Get linear constraint placeholder information
<code>getModelConstraints</code>	Get model constraint placeholder information
<code>getName</code>	Get current name label for optimization function
<code>getNonlcon</code>	Get nonlinear constraint information
<code>getObjectives</code>	Return information about optimization objectives
<code>getObjectivesMode</code>	Return current usage of objective functions
<code>getOperatingPointSets</code>	Return information about optimization operating point sets
<code>getOperatingPointsMode</code>	Return current usage of operating point sets
<code>getParameters</code>	Return information about optimization parameters
<code>getRunInterfaceVersion</code>	Get preferred interface to provide evaluation function
<code>removeConstraint</code>	Remove constraint from optimization
<code>removeFreeVariable</code>	Remove free variable from optimization
<code>removeObjective</code>	Remove objective from optimization
<code>removeOperatingPointSet</code>	Remove operating point set from optimization
<code>removeParameter</code>	Remove parameter from optimization
<code>setConstraintsMode</code>	Set how optimization constraints are to be used
<code>setDescription</code>	Provide description for optimization function
<code>setEnabled</code>	Set enabled status for optimization function
<code>setFreeVariablesMode</code>	Set how optimization free variables are used
<code>setName</code>	Provide name label for optimization function
<code>setObjectivesMode</code>	Set how optimization objective functions are used
<code>setOperatingPointsMode</code>	Set how optimization operating point sets are used
<code>setRunInterfaceVersion</code>	Get preferred interface to provide evaluation function

Methods of `cgoptimstore`

The following methods are available:

evaluate	Evaluate optimization objectives and constraints
evaluateConstraint	Evaluate optimization constraints
evaluateEqCon	Evaluate optimization nonlinear equality constraints
evaluateIneqCon	Evaluate optimization nonlinear inequality constraints
evaluateNonlcon	Evaluate optimization nonlinear constraints
evaluateObjective	Evaluate optimization objectives
get	Get optimization properties
getA	Get linear inequality constraint matrix.
getB	Get linear inequality constraint target values.
getConstraint	Return constraint labels
getDataset	Retrieve data from data set
getFreeVariables	Get optimal values of free variables
getInitFreeVal	Get initial free values for optimization
getLB	Get free variable lower bounds
getLcon	Return linear constraint labels
getNumConstraint	Return number of constraints per label
getNumConstraintLabels	Return number of constraint labels
getNumLcon	Return number of linear constraints per label
getNumLconLabels	Return number of linear constraint labels
getNumNonlcon	Return number of nonlinear constraints per label
getNumNonlconLabels	Return number of nonlinear constraint labels
getNumObjectiveLabels	Return number of objective labels
getNumObjectives	Return number of objectives per label
getNumRowsInDataset	Get number of rows in optimization data set
getObjectives	Return objective labels for optimization
getObjectiveType	Return objective type
getOptimOptions	Retrieve optimization options object
getOutputInfo	Get output information for optimization
getParam	Get optimization parameter
getStopState	Current stop state for optimization
getUB	Get free variable upper bounds
gridEvaluate	Grid evaluation of optimization objectives and constraints
gridPevEvaluate	Grid evaluation of prediction error variance (PEV)
isScalarFreeVariables	Return whether all free variables are scalars
nEvaluate	Natural evaluation of optimization objectives and constraints
nEvaluateConstraint	Natural evaluation of optimization constraints
nEvaluateNonlcon	Natural evaluation of optimization nonlinear constraints
nEvaluateObjective	Natural evaluation of optimization objectives
optimset	Create/alter optimization OPTIONS structure
optimoptions	Create optimization options
pevEvaluate	Evaluate prediction error variance (PEV)
setExitStatus	Set exit status information for optimization
setFreeVariables	Set optimal values of free variables
setOutput	Set diagnostic information for optimization

Functions

addFreeVariable

Add free variable to optimization

Syntax

```
options = addfreeVariable (options, label)
```

Description

A method of `cgoptoptions`. Adds a placeholder for a free variable to the optimization. The string `label` is used to refer to the variable in CAGE.

See Also

Topics

- `setFreeVariablesMode`
- `getFreeVariablesMode`
- `getFreeVariables`
- `removeFreeVariable`

Introduced before R2006a

addLinearConstraint

Add linear constraint to optimization

Syntax

```
options = addLinearConstraint(options, label, A, B)
```

Description

A method of `cgoptimoptions`. Adds a placeholder for a linear constraint to the optimization. The string `label` is used to refer to the constraint in the CAGE GUI. Linear constraints can be written in the form

$$A(1)X(1) + A(2)X(2) + \dots + A(n)X(n) \leq b$$

where $X(i)$ is the i^{th} free variable, A is a vector of coefficients, and b is a scalar bound.

Examples

```
% Add SPK and EGR variables to an optimization
opt = addFreeVariable(opt, 'SPK');
opt = addFreeVariable(opt, 'EGR');
% Add a linear constraint such that 3*SPK - 2*EGR <= 30
opt = addLinearConstraint(opt, 'newCon', [3 -2], 30);
```

See Also

Topics

```
getLinearConstraints
addModelConstraint
setConstraintsMode
removeConstraint
```

Introduced before R2006a

addModelConstraint

Add model constraint to optimization

Syntax

```
options=addModelConstraint(options, label, boundtype, bound)
```

Description

A method of `cgoptimoptions`. Adds a placeholder for a model constraint to the optimization. The string `label` is used to refer to the constraint in CAGE.

`boundtype` can be set either to the string 'greaterthan' or 'lessthan'.

`bound` must be a scalar real.

If `boundtype` = 'greaterthan', the model constraint takes the following form:

CAGE model \geq bound

Similarly, if `boundtype` = 'lessthan', the model constraint takes the form

CAGE model \leq bound

Examples

An optimization requires a constraint where a user-defined function must be less than 500. The following code line adds a placeholder for this constraint that is labeled 'mycon':

```
opt = addModelConstraint(opt, 'mycon', 'lessthan', 500);
```

See Also

Topics

`getModelConstraints`
`addLinearConstraint`
`setConstraintsMode`
`removeConstraint`

Introduced before R2006a

addObjective

Add objective to optimization

Syntax

```
options = addObjective(options, label, typestr)
```

Description

A method of `cgoptimoptions`. Adds a placeholder for an objective function to the optimization. The string `label` is used to refer to the constraint in CAGE.

`typestr` can take one of four values, 'max', 'min', 'min/max', or 'helper'.

Examples

```
opt = addObjective(opt, 'newObj', 'max')
```

Adds an objective function labeled `newObj` to the optimization and indicates that it is to be maximized.

```
opt = addObjective(opt, 'newObj', 'min/max')
```

Adds an objective function labeled `newObj` to the optimization and indicates that the user should be allowed to choose whether it is minimized or maximized from CAGE.

```
opt = addObjective(opt, 'newObj2', 'helper')
```

Adds an objective function labeled `newObj2` to the optimization. The string 'helper' indicates that the function is used as part of the determination of the cost function but is not directly minimized or maximized.

See Also

Topics

- `getObjectives`
- `setObjectivesMode`
- `getObjectivesMode`
- `removeObjective`

Introduced before R2006a

addOperatingPointSet

Add operating point set to optimization

Syntax

```
options = addOperatingPointSet(options, label, vars)
```

Description

A method of `cgoptimoptions`. `options = addOperatingPointSet(options, label, vars)` Adds a placeholder for an additional operating point set to the optimization.

The string `label` is used to refer to the constraint in CAGE. `vars` is a (1-by-N) cell array of strings where $N \geq 1$. Each element of `vars` is a label for a CAGE variable that must appear in the operating point set that the user chooses.

See Also

Topics

```
getOperatingPointSets  
setOperatingPointsMode  
getOperatingPointsMode  
removeOperatingPointSet
```

Introduced before R2006a

addParameter

Add parameter to optimization

Syntax

```
options = addParameter(options, Label, Type, Value)
options = addParameter(options, Label, Type, Value, DisplayName)
```

Description

A method of `cgoptoptions`.

`options = addParameter(options, Label, Type, Value)` adds a parameter to the optimization. The string `Label` is used to refer to the parameter in the `Evaluate` section of your script. You must specify a default value in `Value`. The table below lists the parameter types that are supported along with how to specify their `Type` and `Value`.

Parameter Type	Type	Value
Real number	'number'	Real scalar
Integer	'integer'	Integer scalar
Enumerated list	{'list', {list items}}	One of {list items}
Boolean	'boolean'	true or false

Note: The {list items} cell array for an enumerated list must be a cell array of strings, one for each list member.

You can restrict a numeric parameter ('number' or 'integer') to a valid range. To do this, specify a cell array for `Type` from the following:

Range type	Type
Positive	{TYPESTR, 'positive'}
Negative	{TYPESTR, 'negative'}
User defined	{TYPESTR, [a b]}

where `TYPESTR` is either 'number' or 'integer'. Note that the user-defined range type strictly includes the limits, whereas the positive and negative range types exclude zero. Furthermore, the default `Value` must lie in the specified range.

`options = addParameter(options, Label, Type, Value, DisplayName)` allows you to add a more descriptive label for the parameter in the CAGE Optimization Parameters GUI. Note that you still must refer to the parameter by `label` in the `Evaluate` section of your script.

See Also

Topics

`getParameters`
`getParam`

removeParameter

Introduced before R2006a

evaluate

Evaluate optimization objectives and constraints

Syntax

```
Y = evaluate(optimstore, X)
```

Description

A method of `cgoptimstore`.

Evaluate optimization objectives and constraints.

`Y = evaluate(optimstore, X)` evaluates all of the optimization objectives and constraints at the free variable values `X`. `X` is a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization.

Evaluation over data sets is only supported when the free variables are scalar, that is, you cannot perform evaluation over a data set for "sum" optimizations.

Examples

```
Y = evaluate(optimstore, X, itemnames)
```

evaluates the objectives and constraints specified in the cell array of strings, `itemnames`, at the free variable values `X`. The values of the objectives and constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of objectives and constraints listed in `itemnames`. Note that the evaluation of `Y` is scaled onto [-1 1].

```
Y = evaluate(optimstore, X, itemnames, datasetname)
```

evaluates the specified objectives and constraints at the operating points in the data set specified by the string `datasetname`. `X` must be a (`Nrows`-by-`NfreeVar`) matrix, where `Nrows` is the number of rows in the data set.

```
Y = evaluate(optimstore, X, itemnames, datasetname, rowind)
```

evaluates the specified objectives and constraints at the points of `datasetname` given by `rowind`. `X` must be a (`NRows`-by-`NFreeVar`) matrix where `NRows` is the length of `ROWIND`. `ROWIND` must be a list of integer indices in the range [1 `NumRowsInDataset`]. `Y` is a (`Nrows`-by-`NItems`) matrix.

See Also

Topics

`nEvaluate`
`pevEvaluate`

Introduced before R2006a

evaluateConstraint

Evaluate optimization constraints

Syntax

```
Y = evaluateConstraint(optimstore, X)
```

Description

A method of `cgoptimstore`.

`Y = evaluateConstraint(optimstore, X)` evaluates all of the optimization constraints at the free variable values `X`. `X` must be a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization. The values of the constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of constraints in the optimization.

If you enable scaling of the optimization items, then the evaluation of `Y` is approximately scaled onto `[-1 1]`. See “Scale Optimization” on page 7-49 for more information on scaling.

Negative values of `Y` imply `X` is feasible.

Examples

```
Y = evaluateConstraint(optimstore, X, itemnames)
```

evaluates the constraints specified in the cell array of strings, `itemnames`, at the free variable values `X`. The values of the constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of objectives listed in `itemnames`.

```
[Y, YG] = evaluateConstraint(optimstore, X, itemnames)
```

also evaluates the gradient of the specified constraints in `YG` (if `itemnames` is not specified, then the gradient of all constraints is returned). `YG` is of size `NFreeVar`-by-`NItems`-by-`NPoints`, where `NFreeVar` is the number of free variables in the optimization.

See Also

Topics

`evaluateObjective`
`evaluateNonlcon`

Introduced in R2006b

evaluateEqCon

Evaluate optimization nonlinear equality constraints

Syntax

```
Y = evaluateEqCon(optimstore, X)
```

Description

A method of `cgoptimstore`.

`Y = evaluateEqCon(optimstore, X)` evaluates all of the nonlinear equality constraints in the optimization at the free variable values `X`. `X` must be a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization. The values of the constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of nonlinear equality constraints in the optimization.

If you enable scaling of the optimization items, then the evaluation of `Y` is approximately scaled onto `[-1 1]`. See “Scale Optimization” on page 7-49 for more information on scaling.

Negative values of `Y` imply `X` is feasible.

See Also

Topics

`evaluateIneqCon`

Introduced in R2010b

evaluateIneqCon

Evaluate optimization nonlinear inequality constraints

Syntax

```
Y = evaluateIneqCon(optimstore, X)
```

Description

A method of `cgoptimstore`.

`Y = evaluateIneqCon(optimstore, X)` evaluates all of the nonlinear inequality constraints in the optimization at the free variable values `X`. `X` must be a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization. The values of the constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of nonlinear inequality constraints in the optimization.

If you enable scaling of the optimization items, then the evaluation of `Y` is approximately scaled onto `[-1 1]`. See “Scale Optimization” on page 7-49 for more information on scaling.

Negative values of `Y` imply `X` is feasible.

See Also

Topics

`evaluateEqCon`

Introduced in R2010b

evaluateNonlcon

Evaluate optimization nonlinear constraints

Syntax

```
[varargout] = evaluateNonlcon(optimstore, X, ItemNames)
```

Description

Evaluate optimization nonlinear constraints. A method of `cgoptimstore`.

`Y = evaluateNonlcon(optimstore, X)` evaluates all of the nonlinear constraints in the optimization at the free variable values `X`. `X` must be a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization.

If you enable scaling of the optimization items, then the evaluation of `Y` is approximately scaled onto `[-1 1]`. See “Scale Optimization” on page 7-49 for more information on scaling.

`Y = evaluateNonlcon(optimstore, X, ItemNames)` evaluates the nonlinear constraints specified in the cell array of strings, `ItemNames`, at the free variable values `X`. The values of the nonlinear constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of nonlinear constraints listed in `ItemNames`.

`[Y, YG] = evaluateNonlcon(optimstore, X, ItemNames)` also evaluates the gradient of the specified constraints in `YG` (if `ItemNames` is not specified, then the gradient of all constraints is returned). `YG` is of size `NFreeVar`-by-`NItems`-by-`NPoints`, where `NFreeVar` is the number of free variables in the optimization.

See Also

Topics

`evaluateConstraint`
`evaluateObjective`

Introduced before R2006a

evaluateObjective

Evaluate optimization objectives

Syntax

```
varargout = evaluateObjective(optimstore, X, ItemNames)
```

Description

Evaluate optimization objectives. A method of `cgoptimstore`.

`Y = evaluateObjective(optimstore, X)` evaluates all of the optimization objectives at the free variable values `X`. `X` must be a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization. The values of the objectives are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of objectives in the optimization.

If you enable scaling of the optimization items, then the evaluation of `Y` is approximately scaled onto `[-1 1]`. See “Scale Optimization” on page 7-49 for more information on scaling.

`Y = evaluateObjective(optimstore, X, ItemNames)` evaluates the objectives specified in the cell array of strings, `ItemNames`, at the free variable values `X`. The values of the objectives are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of objectives listed in `ItemNames`.

`[Y, YG] = evaluateObjective(optimstore, X, ItemNames)` also evaluates the gradient of the specified objectives in `YG` (if `ItemNames` is not specified, then the gradient of all objectives is returned). `YG` is of size `NFreeVar`-by-`NItems`-by-`NPoints`, where `NFreeVar` is the number of free variables in the optimization.

See Also

Topics

`evaluateNonlcon`

Introduced before R2006a

get

Get optimization properties

Syntax

```
V = get(optimstore, 'PropertyName')
```

Description

Returns the value of the specified property in the optimization. A method of `cgoptimstore`.

`get(optimstore)` displays all property names and a description of each property for the `OPTIMSTORE` object.

`S = get(optimstore)` returns a structure where each field name is the name of a property of `OPTIMSTORE` and each field contains the description of that property.

Note This method is obsolete. Use the `GETXXX` methods instead.

See Also

Topics

getA
getB

Introduced before R2006a

getA

Get linear inequality constraint matrix.

Syntax

```
A = getA(optimstore)
```

Description

Get the linear inequality constraint matrix. A method of `cgoptimstore`.

`A = getA(optimstore)` returns the linear inequality constraint matrix used in the optimization. `A` is a (NLINCON-by-NFreeVar) matrix where `NFreeVar` is the number of free variables in the optimization and `NLINCON` is the number of linear inequality constraints.

The following code evaluates the linear inequality constraints in the optimization:

```
A = getA(optimstore);  
b = getB(optimstore);  
out = A*x - b;
```

where `x` is a column vector containing the current free variable values.

See Also

Topics

`getB`

Introduced before R2006a

getB

Get linear inequality constraint target values.

Syntax

```
B = getB(optimstore)
```

Description

Get the linear inequality constraint target values. A method of `cgoptimstore`.

`B = getB(optimstore)` returns the linear inequality constraint target values used in the optimization. `B` is a $(\text{NLINCON} - \text{by} - 1)$ column vector where `NLINCON` is the number of linear inequality constraints.

The following code evaluates the linear inequality constraints in the optimization:

```
A = getA(optimstore);  
b = getB(optimstore);  
out = A*x - b;
```

where `x` is a column vector containing the current free variable values.

See Also

Topics

`getA`

Introduced before R2006a

getConstraint

Return constraint labels

Syntax

```
conLabels = getConstraint(optimstore)
```

Description

Return the constraint labels. A method of `cgoptimstore`.

`conLabels = getConstraint(optimstore)` returns the labels for all the constraint functions in optimization. These labels are the those found in the CAGE GUI for the optimization constraints.

See Also

Topics

`getNonlcon`

`getLcon`

Introduced in R2007a

getConstraints

Return information about all optimization constraints

Syntax

```
coninfo = getConstraints(obj)
```

Description

Return information about all optimization constraints. A method of `cgoptimoptions`.

`coninfo = getConstraints(options)` returns a structure array of information regarding the optimization constraint functions. `coninfo(i).label` contains the label for the *i*-th constraint. A string defining the type of the *i*-th constraint is stored in `coninfo(i).typestr`. The constraint parameters are stored in `coninfo(i).pars`.

See Also

Topics

`addModelConstraint`
`addLinearConstraint`

Introduced before R2006a

getConstraintsMode

Return current usage of constraints

Syntax

```
mode = getConstraintsMode(options)
```

Description

Returns a string describing how the optimization makes constraints available to the user. `mode` will be one of `'any'` or `'fixed'`.

See Also

Topics

`setConstraintsMode`

Introduced before R2006a

getDataset

Retrieve data from data set

Syntax

```
V = getDataset(optimstore, datasetName, inputNames)
```

Description

Returns required data from a named data set. A method of `cgoptimstore`.

`PTS = getDataset(optimstore, datasetName)` returns all the data from the specified helper data set. If the data set cannot be found, `data` is returned as empty.

`PTS = getDataset(optimstore, datasetName, inputNames)` returns data from the specified helper data set. Data is retrieved for the columns of the data set with names that match those in `inputNames`. If the dataset cannot be found, `data` is returned as empty.

Examples

```
V = getdataset(optimstore, 'myDS', {'speed', 'afr'})
```

returns a `NPTS` by 2 matrix, `V`.

`NPTS` is the number of rows in the operating point set labeled 'myDS', `V(:, 1)` is the data for the variable labeled 'speed', `V(:, 2)` is the data for the variable labeled 'afr'.

See Also

Topics

`addOperatingPointSet`

Introduced before R2006a

getDescription

Get current description for optimization function

Syntax

```
desc = getDescription(options)
```

Description

A method of `cgoptimoptions`. Returns the description, `desc`, of the user-defined optimization function.

See Also

Topics

`setDescription`

Introduced before R2006a

getEnabled

Get current enabled status for optimization

Syntax

```
en=getEnabled(options)
```

Description

A method of `cgoptimoptions`. Returns whether this user-defined optimization is available to be run. `en` is set to true or false. When an optimization is disabled, the user can still register it with CAGE but is not allowed to create new optimizations using it.

See Also

Topics

`setEnabled`

Introduced before R2006a

getFreeVariables

Get optimal values of free variables

Syntax

```
data = getFreeVariables(obj)
```

Description

A method of `cgoptimstore`. Get the optimal values of the free variables.

`Results = getFreeVariables(obj)` returns the matrix of optimal values that has been set for the free variables. `Results` is a `NSOL` by `NFREEVAR` matrix containing many solutions for the optimal values of the free variables. `NSOL` is the number of solutions and `NFREEVAR` is the number of free variables.

See Also

Topics

`setFreeVariables`

Introduced in R2013a

getFreeVariables

Return optimization free variable labels

Syntax

```
labels=getFreeVariables(options)
```

Description

A method of `cgoptimoptions`. Returns the current placeholder labels for the free variables in the optimization. The labels are returned in a (1-by-NFreeVar) cell array, `labels`, where `NFreeVar` is the number of free variables that have been added to the optimization.

See Also

Topics

`addFreeVariable`
`setFreeVariablesMode`
`getFreeVariablesMode`

Introduced before R2006a

getFreeVariablesMode

Return current usage of free variables

Syntax

```
mode= getFreeVariablesMode(options)
```

Description

A method of `cgoptimoptions`. Returns a string describing how the optimization makes free variables available to the user. `mode` is set to `any` or `fixed`.

See Also

Topics

`setFreeVariablesMode`

Introduced before R2006a

getInitFreeVal

Get initial free values for optimization

Syntax

```
x0 = getInitFreeVal(cos)
```

Description

Get the initial free values for the optimization. A method of `cgoptimstore`.

`x0 = getInitFreeVal(optimstore)` returns the initial values of the free variables used in the optimization. `X0` is a (1-by-`NFreeVar`) matrix where `NFreeVar` is the number of free variables in the optimization.

See Also

Topics

`setFreeVariablesMode`

Introduced before R2006a

getLB

Get free variable lower bounds

Syntax

```
LB = getLB(optimstore)
```

Description

Get the free variable lower bounds. A method of `cgoptimstore`.

`LB = getLB(optimstore)` returns the free variable lower bounds used in the optimization. `LB` is a (1-by-`NFreeVar`) vector where `NFreeVar` is the number of free variables in the optimization.

See Also

Topics

getUB

Introduced before R2006a

getLcon

Return linear constraint labels

Syntax

```
conLabels = getLcon(optimstore)
```

Description

Return the linear constraint labels. A method of `cgoptimstore`.

`conLabels = getLcon(optimstore)` returns the labels for the linear constraints in the optimization. These labels are those found in the CAGE GUI for the optimization linear constraints.

See Also

Topics

`getObjectives`

`getNumNonlcon`

Introduced before R2006a

getLinearConstraints

Get linear constraint placeholder information

Syntax

```
out = getLinearConstraints(options)
```

Description

A method of `cgoptimoptions`. Returns a structure array of information regarding the linear constraints in the optimization. The structure has three fields: `label`, `A`, and `b`. See the help for `addLinearConstraint` for more information on these fields.

See Also

Topics

`addLinearConstraint`
`setConstraintsMode`

Introduced before R2006a

getModelConstraints

Get model constraint placeholder information

Syntax

```
out = getModelConstraints (options)
```

Description

A method of `cgoptimoptions`. Returns a structure array of information regarding the model constraints in the optimization. The structure has three fields: `label`, `boundtype`, and `bound`. See the help for `addModelConstraint` for more information on these fields.

See Also

Topics

`addModelConstraint`
`setConstraintsMode`

Introduced before R2006a

getName

Get current name label for optimization function

Syntax

```
name=getName(options)
```

Description

A method of `cgoptoptions`. Returns the current name label, `name`, for the user-defined optimization function.

See Also

Topics

`setName`

Introduced before R2006a

getNonlcon

Get nonlinear constraint information

Syntax

```
out = getNonlcon(obj)
```

Description

Get nonlinear constraint information. A method of `cgoptimoptions`.

`out = getNonlinearConstraints(options)` returns a structure array of information regarding the nonlinear constraints in the optimization. The structure has three fields: `label`, `type` and `pars`. The `label` field contains the label used for the constraint in the CAGE GUI. The `typestr` field contains constraint type selected by the user. The `pars` field contains any parameters associated with the constraint.

See Also

Topics

`getModelConstraints`
`getLinearConstraints`

Introduced before R2006a

getNumConstraint

Return number of constraints per label

Syntax

```
ncon = getNumConstraint(optimstore)
```

```
ncon = getNumConstraint(optimstore, conLabels)
```

Description

Return the number of constraints per label. A method of `cgoptimstore`.

`ncon = getNumConstraint(optimstore)` returns the number of constraints that will be returned from an evaluation of each labeled constraint. For example, consider an optimization that has a sum constraint over a set of points, S , and a point constraint to be evaluated at each member of S . `NUMCONSTRAINT` will return $[1 \ r]$, where r is the number of points in S .

`ncon = getNumConstraint(optimstore, conLabels)` returns the number of constraints from an evaluation of the defined constraints.

See Also

Topics

`getNumNonlcon`

Introduced in R2007a

getNumConstraintLabels

Return number of constraint labels

Syntax

```
out = getNumConstraintLabels(optimstore)
```

Description

Return the number of constraint labels. A method of `cgoptimstore`.

`out = getNumConstraintLabels(optimstore)` returns the number of constraint labels in the optimization.

See Also

Topics

`getNumObjectiveLabels`

Introduced in R2007a

getNumLcon

Return number of linear constraints per label

Syntax

```
ncon = getNumLcon(optimstore)
```

```
ncon = getNumLcon(optimstore, conLabels)
```

Description

Return the number of linear constraints per label. A method of `cgoptimstore`.

`ncon = getNumLcon(optimstore)` returns the number of constraints that will be returned from an evaluation of each linear constraint.

`ncon = getNumNonlcon(optimstore, conLabels)` returns the number of constraints from an evaluation of the defined constraints.

See Also

Topics

`getNumNonlcon`

`getNumConstraint`

Introduced in R2007a

getNumLconLabels

Return number of linear constraint labels

Syntax

```
numlab = getNumLconLabels(optimstore)
```

Description

Return the number of linear constraint labels. A method of `cgoptimstore`.

`numlab = getNumLconLabels(optimstore)` returns the number of linear constraint labels in the optimization.

See Also

Topics

`getNumConstraintLabels`

Introduced in R2007a

getNumNonlcon

Return number of nonlinear constraints per label

Syntax

```
ncon = getNumNonlcon(optimstore)
```

```
ncon = getNumNonlcon(optimstore, conLabels)
```

Description

Return the number of nonlinear constraints per label. A method of `cgoptimstore`.

`ncon = getNumNonlcon(optimstore)` returns the number of constraints that will be returned from an evaluation of each labeled constraint. For example, consider an optimization that has a sum constraint over a set of points, S , and a point constraint to be evaluated at each member of S . `NCN` will return $[1 \ r]$, where r is the number of points in S .

`ncon = getNumNonlcon(optimstore, conLabels)` returns the number of constraints type for the defined constraints.

See Also

Topics

`getConstraints`

`getNumNonlconLabels`

Introduced before R2006a

getNumNonlconLabels

Return number of nonlinear constraint labels

Syntax

```
numlab = getNumNonlconLabels(optimstore)
```

Description

Returns the number of nonlinear constraint labels in the optimization. A method of `cgoptimstore`.

See Also

Topics

`getNumObjectiveLabels`

Introduced before R2006a

getNumObjectiveLabels

Return number of objective labels

Syntax

```
numlab = getNumObjectiveLabels(optimstore)
```

Description

Returns the number of objective labels in the optimization. A method of `cgoptimstore`.

See Also

Topics

`getNumNonlconLabels`

Introduced before R2006a

getNumObjectives

Return number of objectives per label

Syntax

```
nobj = getNumObjectives(optimstore)
```

```
nobj = getNumObjectives(optimstore, objlabels)
```

Description

Return the number of objectives per label. A method of `cgoptimstore`.

`nobj = getNumObjectives(optimstore)` returns the number of objectives that will be returned from an evaluation of each objective label. For example, consider an optimization that has a sum objective over a set of points, S , and a point objective to be evaluated at each member of S . `nobj` will return $[1\ r]$, where r is the number of points in S .

`nobj = getNumObjectives(optimstore, objlabels)` returns the number of objectives that will be returned for the defined objective labels.

See Also

Topics

`getObjectives`

`getObjectiveType`

Introduced before R2006a

getNumRowsInDataset

Get number of rows in optimization data set

Syntax

```
npts = getNumrowsInDataset(optimstore, datasetName)
```

Description

Returns the number of rows in the named data set. A method of `cgoptimstore`.

Introduced before R2006a

getObjectives

Return objective labels for optimization

Syntax

```
objLabels = getObjectives(optimstore)
```

Description

A method of `cgoptimstore`. Returns the labels for the objective functions in optimization. These labels are those found in the CAGE GUI for the optimization objectives.

See Also

Topics

`getLcon`

Introduced in R2013a

getObjectives

Return information about optimization objectives

Syntax

```
objinfo=getObjectives(options)
```

Description

A method of `cgoptimoptions`. Returns a structure array of information regarding the optimization objective functions. `objinfo(i).label` contains the label for the i^{th} objective. A string defining the type of the i^{th} objective (max, min, min/max, or helper) is stored in `objinfo(i).type`.

See Also

Topics

```
addObjective  
setObjectivesMode  
getObjectivesMode
```

Introduced before R2006a

getObjectivesMode

Return current usage of objective functions

Syntax

```
mode = getObjectivesMode(options)
```

Description

A method of `cgoptimoptions`. Returns a string describing how the optimization makes objectives available to the user. `mode` will be one of `'multiple'`, `'any'`, or `'fixed'`.

See Also

Topics

`setObjectivesMode`

Introduced before R2006a

getOperatingPointSets

Return information about optimization operating point sets

Syntax

```
getOperatingPointSets(options)
```

Description

A method of `cgoptimoptions`. Returns a structure array of information regarding the optimization operating point sets. The structure has two fields, `label` and `vars`. See the help for `addOperatingPointSet` for more information on these fields.

See Also

Topics

```
addOperatingPointSet  
setOperatingPointsMode  
getOperatingPointsMode
```

Introduced before R2006a

getOperatingPointsMode

Return current usage of operating point sets

Syntax

```
mode=getOperatingPointsMode(options)
```

Description

A method of `cgoptoptions`. Returns a string describing how the optimization makes operating point sets available to the user. `mode` will be one of `'default'`, `'fixed'`, or `'any'`.

See Also

Topics

`setOperatingPointsMode`

Introduced before R2006a

getObjectiveType

Return objective type

Syntax

```
objType = getObjectiveType(optimstore)
```

```
objType = getObjectiveType(optimstore, objLabels)
```

Description

Return the objective type. A method of `cgoptimstore`.

`objType = getObjectiveType(optimstore)` returns the objective type of all the objectives in the optimization. A 1-by-NOBJ cell array is returned, each element being 'min', 'max' or 'helper'.

`objType = getObjectiveType(optimstore, objLabels)` returns the objective type for the defined objectives.

See Also

Topics

`getObjectives`

Introduced before R2006a

getOptimOptions

Retrieve optimization options object

Syntax

```
options = getOptimOptions(optimstore)
```

Description

A method of `cgoptimstore`. Returns the optimization configuration object. Information about the optimization set up can be retrieved from this object.

Introduced before R2006a

getOutputInfo

Get output information for optimization

Syntax

```
[exitflag, msg, stats] = getOutputInfo(cos)
```

Description

Get output information for the optimization. A method of `cgoptimstore`.

`[exitflag, termMsg] = getOutputInfo(optimstore)` returns diagnostic output information from `optimstore`. `exitflag` indicates the success (`exitflag > 0`) or failure (`exitflag <= 0`) of the current optimization run. `exitflag` may also give some indication why the optimization terminated. Any termination message set by the optimization can be retrieved from `termMsg`.

`[exitflag, termMsg, output] = getOutputInfo(optimstore)` returns in addition a structure of algorithm-specific information in `output`. For `output` to be non-empty, the user must create it in their algorithm. See the worked example and tutorial for more information on how to create `output` structures.

Introduced before R2006a

getParam

Get optimization parameter

Syntax

```
property_value = getParam(obj, propertyname)
```

Description

Get optimization parameter. A method of `cgoptimstore`.

`V = getParam(optimstore, 'Parameter_name')` returns the value of the specified parameter in the optimization. These optimization parameters must be set up in the Options section of the user-defined script.

See Also

Topics

`addParameter`

“Create a User-Defined Optimization” on page 9-7

Introduced before R2006a

getParameters

Return information about optimization parameters

Syntax

```
getParameters(options)
```

Description

A method of `cgoptimoptions`. Returns a structure array containing information about the parameters that are defined for the optimization. Parameter information is returned in a structure with fields `label`, `typestr`, `value`, and `displayname`. See the help for `addParameter` for more information on these fields.

See Also

Topics

`addParameter`
`getParam`

Introduced before R2006a

getRunInterfaceVersion

Get preferred interface to provide evaluation function

Syntax

```
ver = getRunInterfaceVersion(obj)
```

Description

Get the preferred interface to provide the evaluation function. A method of `cgoptimoptions`.

`ver = getRunInterfaceVersion(options)` returns the Model-Based Calibration Toolbox product Version that is emulated when the optimization function's `evaluate` option is called. If `ver` is set to 2, the interface provided by Model-Based Calibration Toolbox Version 2 software is activated. If `ver` is set to 3, the new interface, which Model-Based Calibration Toolbox Version 3 software defines, is used.

See Also

Topics

`setRunInterfaceVersion`

Introduced before R2006a

getStopState

Current stop state for optimization

Syntax

```
stop= getStopState(opt)
```

Description

A method of `cgoptimstore`. `stop= getStopState(optimstore)` returns the current stop state for the optimization. The stop state could be set by the Stop button on the Running Optimization progress bar or via a call to `setStopState` within a script.

See Also

Topics

`setStopState`

Introduced before R2006a

getUB

Get free variable upper bounds

Syntax

```
UB = getUB(optimstore)
```

Description

A method of `cgoptimstore`. Returns the free variable upper bounds used in the optimization. `UB` is a (1-by-`NFreeVar`) vector where `NFreeVar` is the number of free variables in the optimization.

See Also

Topics

`getLB`

Introduced before R2006a

gridEvaluate

Grid evaluation of optimization objectives and constraints

Syntax

```
Y = gridEvaluate(optimstore, X)
```

```
Y = gridEvaluate(optimstore, X, objconname)
```

```
Y = gridEvaluate(optimstore, X, objconname, datasetname)
```

```
Y = gridEvaluate(optimstore, X, objconname, datasetname, rowind)
```

Description

A method of `cgoptimstore`.

`Y = gridEvaluate(optimstore, X)` evaluates all the objectives and constraints at the points `X` for the current run. This call produces identical results to the equivalent call to `cgoptimstore/evaluate`.

`Y = gridEvaluate(optimstore, X, objconname)` evaluates the objectives/constraints specified in the cell array `objconname` as described above.

`Y = gridEvaluate(optimstore, X, objconname, datasetname)` evaluates all the objectives and constraints at all combinations of the points in `datasetname` with `X`. The return matrix, `Y`, is of size `SIZE(X,1) - by - (NOBJ+NCON) - by - NPTS`, where `NOBJ` is the number of objectives, `NCON` is the number of constraints and `NPTS` is the number of rows in `P`. Further, `Y(I, J, K)` is the value of the `J`-th objective/constraint at `X(I, :)` and `P(K, :)`. `Y` is scaled on `[-1 1]`.

Examples

Objectives : 01, 02

Constraints : C1, C2

Primary data set:

A	B
4	5
1	3

Free variables:

X1	X2	X3
2	4	8

X1	X2	X3
1	9	3
6	2	7

X

In this case, the following command

```
Y = gridEvaluate(optimstore, X)
```

evaluates objectives and constraints at the following points:

A	B	X1	X2	X3
4	5	2	4	8
4	5	1	9	3
4	5	6	2	7
1	3	2	4	8
1	3	1	9	3
1	3	6	2	7

Y is a 3-by-4-by-2 matrix where

Y(:, 1, 1) = Values of O1 at A = 4, B = 5

Y(:, 2, 1) = Values of O2 at A = 4, B = 5

Y(:, 3, 1) = Values of C1 at A = 4, B = 5

Y(:, 4, 1) = Values of C2 at A = 4, B = 5

Y(:, 1, 2) = Values of O1 at A = 1, B = 3

Y(:, 2, 2) = Values of O2 at A = 1, B = 3

Y(:, 3, 2) = Values of C1 at A = 1, B = 3

Y(:, 4, 2) = Values of C2 at A = 1, B = 3

```
Y = gridEvaluate(optimstore, X, objconname, datasetname, rowind)
```

evaluates the specified objectives/constraints at the points of datasetname given by rowind as described above. Y is a length(rowind) by length(objconname) by npts matrix.

See Also

Topics

evaluate

Introduced before R2006a

gridPevEvaluate

Grid evaluation of prediction error variance (PEV)

Syntax

```
[y, ysums] = gridpevevaluate(optimstore, X)
```

```
Y = gridpevevaluate(optimstore, X, objconname)
```

```
Y = gridpevevaluate(optimstore, X, objconname, datasetname)
```

```
Y = gridpevevaluate(optimstore, X, objconname, datasetname, rowind)
```

Description

Warning The evaluation of PEV is no longer supported in `cgoptimstore` and this method will return PEV values of zero (as detailed below) if called.

A method of `cgoptimstore`.

`Y = gridpevevaluate(optimstore, X)` produces identical results to the equivalent call to `cgoptimstore/pevEvaluate`

`Y = gridpevevaluate(optimstore, X, objconname)` returns PEV values of zero for the objectives/constraints specified in the cell array `objconname`.

`Y = gridpevevaluate(optimstore, X, objconname, datasetname)` returns PEV values of zero for the specified objectives/constraints. The return matrix, `Y`, is of size `SIZE(X,1) - by - (NOBJCON) - by - NPTS`, where `NOBJCON` is the number of specified objectives/constraints and `NPTS` is the number of rows in `P`.

`Y = gridpevevaluate(optimstore, X, objconname, datasetname, rowind)` returns PEV values of zero for the specified objectives/constraints. `Y` is a `LENGTH(ROWIND) by LENGTH(OBJCONNAME) by NPTS` matrix.

See Also

Topics

`pevEvaluate`

Introduced before R2006a

isScalarFreeVariables

Return whether all free variables are scalars

Syntax

```
stat = isScalarFreeVariables(optimstore)
```

Description

Return whether all the free variables are scalars. A method of `cgoptimstore`.

`stat = isScalarFreeVariables(optimstore)` returns `TRUE` if all the free variables are scalars and `FALSE` otherwise.

Introduced before R2006a

nEvaluate

Natural evaluation of optimization objectives and constraints

Syntax

```
[y, ysums] = nEvaluate(optimstore, x)
```

```
Y = nEvaluate(optimstore, x, itemNames)
```

```
Y = nEvaluate(optimstore, x, itemNames, datasetName)
```

```
Y = nEvaluate(optimstore, x, itemNames, datasetName, rowind)
```

Description

Natural evaluation of optimization objectives and constraints. A method of `cgoptimstore`.

`Y = nEvaluate(optimstore, x)` evaluates the raw values of all of the optimization objectives and constraints at the free variable values `X`. `X` is a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization.

`Y = nEvaluate(optimstore, x, itemNames)` evaluates the raw values of the objectives and constraints specified in the cell array of strings, `itemNames`, at the free variable values `X`. The values of the objectives and constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of objectives and constraints listed in `itemNames`.

`Y = nEvaluate(optimstore, x, itemNames, datasetName)` evaluates the specified objectives and constraints at the operating points in the data set specified by the string `datasetName`.

`Y = nEvaluate(optimstore, x, itemNames, datasetName, rowind)` evaluates the specified objectives and constraints at the points of `datasetName` given by `rowind`. `X` must be a (`NRows`-by-`NFreeVar`) matrix where `NRows` is the length of `rowind`. `rowind` must be a list of integer indices in the range `[1 NumRowsInDataset]`. `Y` is a (`NRows`-by-`NItems`) matrix.

See Also

Topics

evaluate

Introduced before R2006a

nEvaluateConstraint

Natural evaluation of optimization constraints

Syntax

```
Y = nEvaluateConstraint(optimstore, x)
```

```
Y = nEvaluateConstraint(optimstore, x, itemNames)
```

Description

A method of `cgoptimstore`.

`Y = nEvaluateConstraint(optimstore, X)` evaluates all of the optimization constraints at the free variable values `x`. `X` must be a (NPoints-by-NFreeVar) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization. The raw values of the constraints are returned in `Y`, which is of size (NPoints-by-NItems) where `NItems` is the number of constraints in the optimization.

`Y = nEvaluateConstraint(optimstore, X, itemNames)` evaluates the constraints specified in the cell array of strings, `itemNames`, at the free variable values `X`. The raw values of the constraints are returned in `Y`, which is of size (NPoints-by-NItems) where `NItems` is the number of constraints listed in `itemNames`.

See Also

Topics

`evaluateObjective`

`evaluateNonlcon`

Introduced in R2006b

nEvaluateNonlcon

Natural evaluation of optimization nonlinear constraints

Syntax

```
y = nEvaluateNonlcon(optimstore, x)
```

```
Y = nEvaluateNonlcon(optimstore, x, itemNames)
```

Description

Natural evaluation of optimization nonlinear constraints. A method of `cgoptimstore`.

`Y = nEvaluateNonlcon(optimstore, x)` evaluates all of the optimization nonlinear constraints at the free variable values `X`. `X` must be a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization. The raw values of the constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of nonlinear constraints in the optimization.

`Y = nEvaluateNonlcon(optimstore, x, itemNames)` evaluates the nonlinear constraints specified in the cell array of strings, `itemNames`, at the free variable values `X`. The raw values of the constraints are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of nonlinear constraints listed in `itemNames`.

See Also

Topics

`evaluateObjective`

`evaluateNonlcon`

Introduced before R2006a

nEvaluateObjective

Natural evaluation of optimization objectives

Syntax

```
y = nEvaluateObjective(optimstore, x)
```

```
Y = nEvaluateObjective(optimstore, x, itemNames)
```

Description

Natural evaluation of optimization objectives. A method of `cgoptimstore`.

`Y = nEvaluateObjective(optimstore, x)` evaluates all of the optimization objectives at the free variable values `X`. `X` must be a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization. The raw values of the objectives are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of objectives in the optimization.

`Y = nEvaluateObjective(optimstore, x, itemNames)` evaluates the objectives specified in the cell array of strings, `itemNames`, at the free variable values `X`. The raw values of the objectives are returned in `Y`, which is of size (`NPoints`-by-`NItems`) where `NItems` is the number of objectives listed in `itemNames`.

See Also

Topics

`evaluateObjective`

`evaluateNonlcon`

Introduced before R2006a

optimset

Create/alter optimization OPTIONS structure

Syntax

```
options = optimset(optimstore)
options = optimset(olddopts, optimstore)
options = optimset(optimfunction, optimstore)
options = optimset(..., 'param1',value1,...)
```

Description

Create/alter optimization OPTIONS structure. A method of `cgoptimstore`.

`options = optimset(optimstore)` creates an optimization options structure that can be used with Optimization Toolbox functions, with the named parameters altered with the specified values. Any parameters specified in the optimization that match (by name) those in the default options structure are copied into `options`.

`options = optimset(olddopts, optimstore)` creates a copy of `olddopts` and copies matching parameters from the optimization into it.

`options = optimset(optimfunction, optimstore)` creates an options structure with all the parameter names and default values relevant to the optimization function named in `optimfunction` and then copies matching parameters from the optimization into it.

`options = optimset(..., 'param1',value1,...)` sets the additional named parameters to the specified values.

See Also

`getParam` | `optimoptions`

Topics

“Choose Between `optimoptions` and `optimset`”

Introduced before R2006a

optimoptions

Create optimization options

Syntax

```
options = optimoptions(optimstore)
options = optimset(olddopts, optimstore)
options = optimoptions(optimfunction, optimstore)
options = optimoptions(..., 'param1',value1,...)
```

Description

Create/alter optimization OPTIONS structure. Method of `cgoptimstore`.

`options = optimoptions(optimstore)` creates an optimization options structure that can be used with Optimization Toolbox functions. Any parameters specified in the optimization that match (by name) those in the default options structure are copied into `options`.

`options = optimoptions(olddopts, optimstore)` creates a copy of `olddopts` and copies matching parameters from the optimization into it.

`options = optimoptions(optimfunction, optimstore)` creates an options structure with all the parameter names and default values relevant to the optimization function named in `optimfunction` and then copies matching parameters from the optimization into it.

`options = optimset(..., 'param1',value1,...)` sets the additional named parameters to the specified values.

See Also

`getParam` | `optimset`

Topics

“Choose Between `optimoptions` and `optimset`”

Introduced in R2021a

pevEvaluate

Evaluate prediction error variance (PEV)

Syntax

```
Y = pevEvaluate(optimstore, X)
```

Description

Warning The evaluation of PEV is no longer supported in `cgoptimstore` and this method will return PEV values of zero (as detailed below) if called.

A method of `cgoptimstore`.

```
Y = pevEvaluate(optimstore, X, itemnames)
```

returns PEV values of zero for objectives/constraints at the free variable values `X`. `X` is a (`NPoints`-by-`NFreeVar`) matrix where `NPoints` is the number of points to be evaluated and `NFreeVar` is the number of free variables in the optimization.

```
Y = pevEvaluate(optimstore, X, objconname, datasetname)
```

returns PEV values of zero for the objectives/constraints at the operating points in the data set specified by the string `datasetname`.

```
Y = pevEvaluate(optimstore, X, objconname, datasetname, rowind)
```

returns PEV values of zero for the specified objectives/constraints at the points of `datasetname` given by `rowind`. `X` must be a (`NRows`-by-`NFreeVar`) matrix where `NRows` is the length of `rowind`. `rowind` must be a list of integer indices in the range `[1 NumRowsInDataset]`. `Y` is a (`NRows`-by-`NItems`) matrix.

See Also

Topics

`gridPevEvaluate`

Introduced before R2006a

removeConstraint

Remove constraint from optimization

Syntax

```
obj = removeConstraint(obj, sLabel)
```

Description

Remove a constraint from the optimization. A method of `cgoptimoptions`.

`obj = removeConstraint(options, label)` removes the placeholder for the constraint referred to by the string `label`.

See Also

Topics

- `getModelConstraints`
- `getLinearConstraints`
- `addModelConstraint`
- `addLinearConstraint`

Introduced before R2006a

removeFreeVariable

Remove free variable from optimization

Syntax

```
obj = removeFreeVariable(obj, sLabel)
```

Description

Remove a free variable from the optimization. A method of `cgoptimoptions`.

`options = removeFreeVariable(options, label)` removes the placeholder for the free variable referred to by the string `label`.

See Also

Topics

`getFreeVariables`
`addFreeVariable`

Introduced before R2006a

removeObjective

Remove objective from optimization

Syntax

```
obj = removeObjective(obj, sLabel)
```

Description

Remove an objective from the optimization. A method of `cgoptimoptions`.

`options = removeObjective(options, label)` removes the placeholder for the objective referred to by the string `label`.

See Also

Topics

`getObjectives`

`addObjective`

Introduced before R2006a

removeOperatingPointSet

Remove operating point set from optimization

Syntax

```
obj = removeOperatingPointSet(obj, sLabel)
```

Description

Remove an operating point set from the optimization. A method of `cgoptimoptions`.

`options = removeOperatingPointSet(options, label)` removes the placeholder for the operating point set referred to by the string `label`.

See Also

Topics

`getOperatingPointSets`

`addOperatingPointSet`

Introduced before R2006a

removeParameter

Remove parameter from optimization

Syntax

```
obj = removeParameter(obj, sLabel)
```

Description

Remove a parameter from the optimization. A method of `cgoptimoptions`.

Removes the placeholder for the parameter referred to by the string `label`.

See Also

Topics

`getParameters`
`addParameter`

Introduced before R2006a

setConstraintsMode

Set how optimization constraints are to be used

Syntax

```
options=setConstraintsMode(options, modestr)
```

Description

A method of `cgoptoptions`. Sets the mode that governs how the user can set up constraints for the optimization in CAGE.

When `modestr = any`, the user can add any number of constraints.

When `modestr = fixed`, the user can only edit the constraints that are added by the user-defined optimization function.

See Also

Topics

`getConstraintsMode`
`addModelConstraint`
`addLinearConstraint`

Introduced before R2006a

setDescription

Provide description for optimization function

Syntax

```
options=setDescription(options, desc)
```

Description

A method of `cgoptoptions`. Sets the description for the optimization object to be the string `desc`.

See Also

Topics

`getDescription`

Introduced before R2006a

setEnabled

Set enabled status for optimization function

Syntax

```
options = setEnabled(options, status)
```

Description

A method of `cgoptoptions`. Sets the optimization function enabled status. `status` must be true or false. When an optimization is disabled, you can still register it with CAGE but are not allowed to create new optimizations using it.

See Also

Topics

`setEnabled`

Introduced before R2006a

setExitStatus

Set exit status information for optimization

Syntax

```
optimstore = setExitStatus(optimstore, exitflag, termmsg)
```

Description

Set exit status information for the optimization. A method of `cgoptimstore`.

`optimstore = setExitStatus(optimstore, exitflag, termmsg)` sets termination status information in the `optimstore`. `exitflag` is an integer which determines whether the optimization has terminated successfully. A value of `exitflag > 0` indicates success, and `exitflag <= 0` indicates failure. In any event, a termination message can be passed back to the optimization through `termmsg`.

See Also

Topics

“Create a User-Defined Optimization” on page 9-7

Introduced before R2006a

setFreeVariables

Set optimal values of free variables

Syntax

```
OUT = setFreeVariables(optimstore, results)
```

Description

Sets the optimal values of the free variables, as returned by the optimization, into the `optimstore`. A method of `cgoptimstore`.

`results` is a `npts` by `nfreevar` matrix containing the optimal values of the free variables. `nsol` is the number of solutions and `nfreevar` is the number of free variables.

Note This function *must* be called at the end of the optimization for the optimal values to be stored.

See Also

Topics

`getFreeVariables`

Introduced before R2006a

setFreeVariablesMode

Set how optimization free variables are used

Syntax

```
options = setFreeVariablesMode(options, modestr)
```

Description

A method of `cgoptoptions`. Sets the mode that governs how the user is allowed to set up free variables for the optimization in the CAGE GUI.

When `modestr = 'any'`, the user is allowed to add any number of free variables.

When `modestr = 'fixed'`, the user is only allowed to use the number of free variables that are added by the user-defined optimization function.

See Also

Topics

`getFreeVariablesMode`
`addFreeVariable`

Introduced before R2006a

setName

Provide name label for optimization function

Syntax

```
options = setName(options, name)
```

Description

A method of `cgoptimoptions`. Sets the name label for the optimization object to be the string name.

See Also

Topics

`getName`

Introduced before R2006a

setObjectivesMode

Set how optimization objective functions are used

Syntax

```
options = setObjectivesMode(options, modestr)
```

Description

A method of `cgoptimoptions`. Sets the mode that governs whether the user is allowed to set up objectives for the optimization in the CAGE GUI.

When `modestr = 'any'`, the user is allowed to add any number of objectives.

When `modestr = 'fixed'`, the user is only allowed to edit the objectives that are added by the user-defined optimization function.

When `modestr = 'multiple'`, the user is only allowed to run the optimization if he or she has defined two or more objectives.

See Also

Topics

`getObjectivesMode`
`addObjective`

Introduced before R2006a

setOperatingPointsMode

Set how optimization operating point sets are used

Syntax

```
options = setOperatingPointsMode(options, modestr)
```

Description

A method of `cgoptimoptions`. Sets the mode that governs how the user is allowed to set up operating point sets for the optimization in CAGE.

When `modestr = 'any'`, the user is allowed to add any number of operating point sets.

When `modestr = 'default'`, the user is allowed to optionally define a single operating point set to run the optimization over.

When `modestr = 'fixed'`, the number of operating point sets required can be fixed by the optimization function and the user is not allowed to add or remove any using the CAGE GUI.

See Also

Topics

`getOperatingPointsMode`
`addOperatingPointSet`

Introduced before R2006a

setOutput

Set diagnostic information for optimization

Syntax

```
optimstore = setOutput(optimstore, OUTPUT)
```

Description

Set diagnostic information for the optimization. A method of `cgoptimstore`.

`optimstore = setOutput(optimstore, OUTPUT)` sets diagnostic information for the optimization in `optimstore`. Any diagnostic information is passed to `optimstore` through the structure, `OUTPUT`. See the worked example for an example of creating an `OUTPUT` structure.

See Also

Topics

“Create a User-Defined Optimization” on page 9-7

Introduced before R2006a

setOutputInfo

Set output information for optimization

Syntax

```
optimstore = setOutputInfo (optimstore, exitflag, termmsg, output)
```

Description

Sets output information for the optimization in `optimstore`. A method of `cgoptimstore`.

The following information is set:

- `exitflag`: integer value status flag indicating why the optimization has terminated. `exitflag > 0` implies that the optimization has terminated successfully.
- `termmsg`: Message that is displayed at termination of algorithm. Normally used for error messages.
- `output`: Structure of algorithm statistics for the optimization.

Note This method is obsolete. Use `cgoptimstore/setExitStatus` and `cgoptimstore/setOutput` instead.

See Also

Topics

`setExitStatus`

`setOutput`

Introduced before R2006a

setRunInterfaceVersion

Get preferred interface to provide evaluation function

Syntax

```
obj = setRunInterfaceVersion(obj, ver)
```

Description

Set the preferred interface to provide the evaluation function. A method of `cgoptimoptions`.

Sets the Model-Based Calibration Toolbox product Version that is emulated when the optimization function's `evaluate` option is called. If `ver` is set to 2, the interface provided by Model-Based Calibration Toolbox Version 2 software is activated. If `ver` is set to 3, the new interface, which Model-Based Calibration Toolbox Version 3 software defines, will be used.

The interface version that the current version of the Model-Based Calibration Toolbox product runs is superior in its capabilities, however it does contain some backwards incompatibilities with the interface used in version 2. You can use this function in old Model-Based Calibration Toolbox optimization files that fail to work with the newer interface.

See Also

Topics

`getRunInterfaceVersion`

Introduced before R2006a

setStopState

Set current stop state for optimization

Syntax

```
setStopState(opt, stop)
```

Description

Set current stop state for optimization. A method of `cgoptimstore`.

`stop = setStopState(optimstore, stop)` sets the current stop state (TRUE or FALSE) for the optimization. Note that this command does not stop an optimization, the optimization script must do this.

See Also

Topics

`getStopState`

Introduced before R2006a

Data Sets

This section includes the following topics:

- “Use Data Sets Views” on page 10-2
- “Set Up Data Sets” on page 10-4
- “View Data in a Table” on page 10-11
- “Plot Outputs” on page 10-12
- “Use Color to Display Information” on page 10-14
- “Link Factors in a Data Set” on page 10-17
- “Manipulate Models in Data Set View” on page 10-19
- “Export Data Sets” on page 10-20

Use Data Sets Views




You can use the **Data Set** view for these main functions:

- Validating calibrations with experimental data
- Filling tables by reference to a set of experimental data
- Constructing operating point sets for running optimizations
- Investigating optimization results and using them to fill tables

For worked examples about data sets, see:

Data Sets consists of four views. These views display different aspects of the data set. Each view is accessible from the **View** menu or by clicking the appropriate button on the toolbar.



- **Factor Information**

List of all available project expressions, which can be added to the data set for display and evaluation.

- **View Data**


Displays the data in a table. Individual entries can be altered. Columns of data can be assigned to CAGE expressions.

- **Plot Outputs**

Displays models and features evaluated at the data points (of the data set).

- **Fill Table from Data Set**

This mode allows you to fill lookup tables by reference to experimental data.



Set Up Data Sets

In this section...

“How to Set Up Data Sets” on page 10-4

“Importing Experimental Data from File” on page 10-4

“Importing Data from the Model Browser” on page 10-6


“Importing Data from a Table in Your Session” on page 10-7

“Merging Data Sets” on page 10-7

“Specifying the Factors Manually” on page 10-7

“Creating a Factor from the Error Between Factors” on page 10-9

How to Set Up Data Sets

The **Data Sets** view displays the strategies, tables, and models, etc., as a list of factors in the default **Data Set Factors** view. You can also display the same factors as columns in a grid, with all factors displayed as columns in the list, by selecting the View Data toolbar button (). The data set works over a grid of values, which is not necessarily the same as the normalizers of any included tables in the data set.

You have to set the input factors and their values to define the grid in the data set. You can do this in one of these ways:

- Import experimental data from file. See “Importing Experimental Data from File” on page 10-4.
- Import data from the Model Browser. See “Importing Data from the Model Browser” on page 10-6.
- Import the values from a table in your CAGE session. See “Importing Data from a Table in Your Session” on page 10-7.
- Merge data sets that share the same factors. See “Merging Data Sets” on page 10-7.
- Specify the factors and their values manually. See “Specifying the Factors Manually” on page 10-7.

The next sections describe each of these in detail.

Importing Experimental Data from File

You can import experimental data to a data set, either to validate a calibration or to use it as the basis for a calibration.

You can import data that is stored in the following formats:

- Microsoft Excel spreadsheets
- Comma-separated value files
- MAT-files
- Data in the Model Browser

Data Format for Importing from Excel or Comma-Separated Value

When you import data from either a Microsoft Excel spreadsheet or from a comma-separated value file, you must ensure that the data is organized in the following manner:

- The first row can either be column headers (text) or entries (numbers).
- The second row can be a row of units (text), as for importing into the Model Browser. CAGE ignores this row.
- All the other row and column entries must be numbers.

Note The Data Editor can create a tailor-made Excel sheet for you to fill with data and then import. This sheet will be in the format the Model-Based Calibration Toolbox product expects to import data. See “Tailor-Made Excel Sheets” in the *Model-Based Calibration Toolbox Model Browser User's Guide*.

Importing from MAT-files

When you import from a MAT-file, you must ensure that the file contains numbers only, that is, a double array.

To import experimental data:

- 1 Select **File > New > Data Set** to create an empty data set.
- 2 Select **File > Import > Data > File**.
- 3 In the file browser, select the correct file to import. This opens the Loading Data from *MAT-filename* dialog. Use this dialog to select the data in the MAT file you wish to import into the data set. Click **OK**.

This opens the **Data Set Import Wizard**.

- 4 Discard any columns of data you do not want to import by selecting the column and clicking the button shown.




- 5 Click **Next**.

The following screen asks you to associate variables in your project with data columns in the data.

- 6 Highlight the variable in the **Project Assignments** column and the corresponding data column in the **Data Column**, then click the assign button, shown.



- 7 Repeat step 5 until you are satisfied that you have associated all the variables and data columns. Any unassigned data columns are treated as output factors.



- 8 Click **Finish** to close the dialog box.

CAGE imports your data and you can view your data set.

Importing Data from the Model Browser

You can import data sets from a project currently loaded in the Model Browser.

To import data from the Model Browser:

- 1 Select **File > New > Data Set** to create an empty data set.
- 2 Select **File > Import > Data > Model Browser**.
- 3 In the dialog box, select the correct data set in the current Model Browser project to import.

The Data Set Import Wizard opens.

- 4 Select the data columns that you want to import into the data set. Exclude any columns of data you do not want to import. To do so, select the column and click the button shown.



If your project is empty, you can click **OK** to import your selections. Otherwise, click **Next**.

- 5 On the next screen, highlight variables in the **Project Assignments** column to associate variables in your project with data columns in the data. Then, highlight the corresponding data column under **Data Column**, and click the assign button.
- 6 Click **Finish** to close the wizard and import the data.

See also “Export Data Sets” on page 10-20.

Importing Data from a Table in Your Session

To import data from a table:

- 1 Select **Data > Import > Import from Table**.

If your data set already contains data, a dialog box asks whether you want to **Fill** the data set from the table or **Overwrite** the data set from the table.

- Select **Fill** to use the table values to fill the factors in your data set.
- Select **Overwrite** to disregard all factors in your data set and fill the data set with the input and output factors from the table.

A dialog box opens.

- 2 Select the correct table from your session to import, and click **OK**.

When you have imported your data, you can view the data set.

Merging Data Sets

To merge another data set in your project with the currently selected data set:

- 1 Select **Data -> Import -> Merge Data Set**.

The Merge Data Sets dialog box appears containing a list of all data sets in your project.

- 2 Select the data set you want to merge with the current data set, and click **OK**.

Columns of inputs and external data are appended to columns with names that match in the current data set.

Outputs (models) and any other columns without matching names are not merged.

The values for any unmatched columns are set to the set point if possible, or zero otherwise.

Specifying the Factors Manually


- 1 Select the **Data Set** view by clicking the large **Data Sets** button in the **Data Objects** pane.
- 2 Add a data set to the project by selecting **File -> New -> Data Set**.
- 3 Select the factors. (See “Selecting the Factors” on page 10-7.)
- 4 Build the grid. (See “Manually Setting Values of the Input Variables” on page 10-9.)

After you complete these steps you can view the data set.


This section describes

- “Selecting the Factors” on page 10-7
- “Manually Setting Values of the Input Variables” on page 10-9

Selecting the Factors

Clicking the Factors View button in the toolbar (). This displays two list boxes.

- The upper list shows all factors within the data set. You can sort factors by clicking the column headings.
- The lower list shows CAGE project expressions.



You can use this view to add factors to or remove factors from the data set.

To add a factor to a data set,

- Right-click a factor and select **Add to Data Set** from the context menu.
- Alternatively, select the factor or factors that you want to add to the data set from the list in the lower **Project Expressions** pane, then select **Data > Factors > Add to Data Set**.


To make multiple selections, use the standard **Shift+click** or **Ctrl+click**.

To remove a factor from a data set,

- 1 Select the factor or factors that you want to remove from the data set.
- 2 Right-click and select **Remove from Data Set**, or select the menu item **Data -> Factors -> Remove From Data Set**.

Note Links between the two lists are always preserved, so clicking load in the upper list also selects load in the lower list. In other words, you can copy or remove from either list and the relevant results appear in both.

Manually Setting Values of the Input Variables

Clicking the Build Grid toolbar button () or selecting **Data -> Build Grid** enables you to set the values of the input variables for the data set.

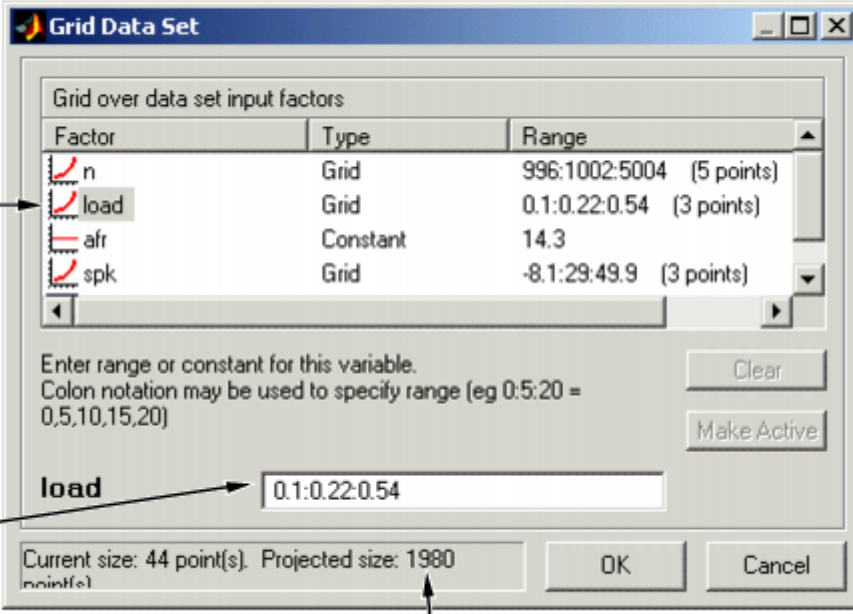
To build a full factorial grid,

- 1 Select **Data -> Build Grid**.
- 2 Select the factor that you want to define a grid for.
- 3 Set the grid for the factor.

To set a grid of 5, 10, 15, 20, 25, 30, input the following: 5:5:30, where the first number is the minimum, the second is the step size, and the last number is the maximum value.

- 4 Check the size of the data set in the pane. The current size reported at the bottom of the dialog is the size if you click **Cancel** to leave the data set unchanged. The projected size is created if you click **OK**. In the following example, the projected size of 45 you can see is obtained by multiplying the number of points for each factor with a grid (in this case, $3 * 5 * 3$).
- 5 Select the next factor that you want to define a grid for.
- 6 When you have set the grids for all the factors, click **OK**.

1. Highlight the input factor.



Factor	Type	Range
n	Grid	996:1002:5004 (5 points)
load	Grid	0.1:0.22:0.54 (3 points)
afr	Constant	14.3
spk	Grid	-8.1:29:49.9 (3 points)

Enter range or constant for this variable.
Colon notation may be used to specify range (eg 0:5:20 = 0,5,10,15,20)

load 0.1:0.22:0.54

Current size: 44 point(s). Projected size: 1980

2. Set the range for the factor. 3. Check the size of the data set.


Creating a Factor from the Error Between Factors

To create a factor that is the difference between two other factors,


- 1** Highlight the two factors, using **Ctrl+click** or **Shift+click**.
- 2** Select **Create Error** from the right-click menu on either column head.

This creates a new factor that is the difference between the two other factors.







View Data in a Table

Click the **View Data** button () in the toolbar or select **View -> Data** to display the data in tabular form and a list of the current items in the project.

Note that this view is only enabled if you have a grid of points at which to evaluate and display the models and variables. This grid is not necessarily derived from the normalizers of any tables included in the data set. You can set the grid by importing experimental or table data, or by using the Build

Grid toolbar button (). See “Set Up Data Sets” on page 10-4.

Inputs to the selected column, colored cream Input that is not an input to the selected column Selected column

	 n	 load	 afr	 spk	rmeas	tgmeas	 Torque: Model	 Torque: Strategy
1	2235	0.549	9.5	0.1	2247	66.7	71.666	66.079
2	3591	0.454	13.2	0.1	3613	54.1	47.163	46.891
3	4946	0.651	12	0.1	4974	73.7	47.573	79.256
4	881	0.648	11.9	5.7	881	75.8	99.23	80.211
5	2234	0.441	13.3	0.1	2247	55.9	51.256	45.152
6	3591	0.747	10.9	0.1	3612	90	92.837	105.586
7	4947	0.541	9.7	0.1	4973	62.8	57.76	57.587
8	881	0.622	9.9	0.1	884	72.1	76.198	60.926
9	1219	0.333	14	0.1	1224	41.8	33.226	21.318
10	1558	0.382	12	0.1	1567	49.4	40.487	31.957
11	1896	0.209	10.7	3.3	1906	28.5	3.492	4.197
12	2234	0.284	9.8	3.2	2245	36	23.063	19.891
13	2574	0.407	13.4	3	2588	49.9	49.629	44.794
14	2914	0.595	11.5	3.1	2929	70.5	84.68	82.229
15	3251	0.781	12.3	3.1	3268	90.5	117.424	117.259
16	3589	0.668	13.5	3	3608	77.1	87.987	96.408
17	3930	0.452	11.9	3.1	3952	52.7	46.511	51.722
18	4268	0.235	10.9	3	4293	27.7	5.253	3.085
19	4606	0.194	12	3.2	4633	21.3	-2.088	-5.771

Columns are color coded by factor type:

- Input factors are white.
- Output factors are gray.

Selecting an output column highlights the input columns associated with it by turning the header cells cream.

Standard editing facilities are available. Double-click an input cell to edit the value.


Cut and paste using the desktop clipboard. Cells, columns, and rows can be copied directly to and from other applications (for example, Excel).

Note You can only edit input values, not output values.

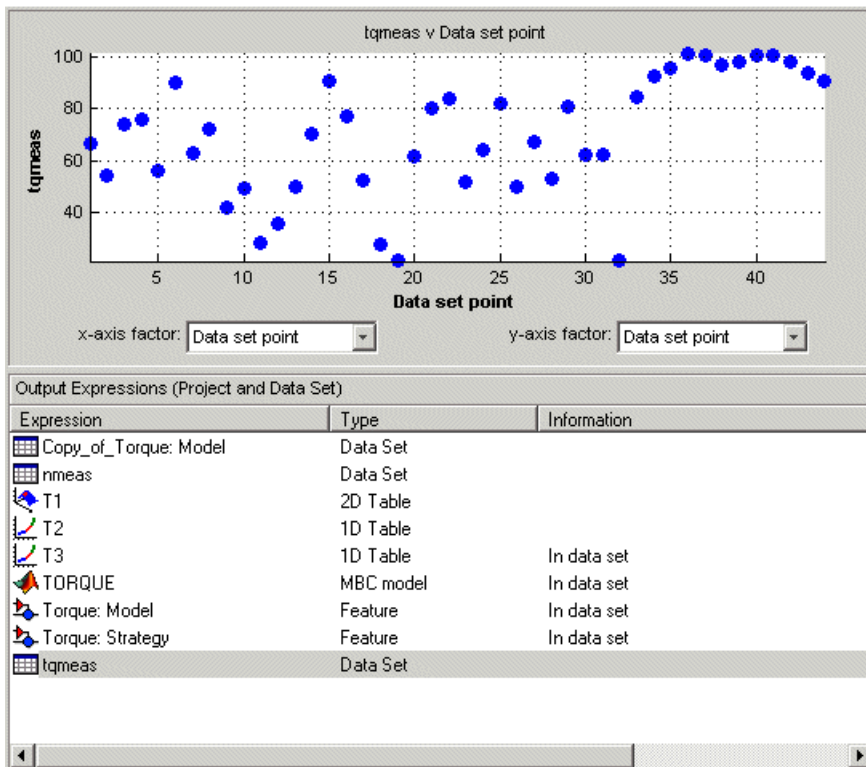
Plot Outputs

Use this to plot the outputs of your data sets.

To view a plot,

- 1 Select **View > Plot** or click the  toolbar button.
- 2 Select an expression from the list to view.

A plot of the selected output factor appears in the top pane.



- 3 Use the pop-up menus below the plot to change the factors displayed.

To zoom in on an area of interest,


- Press both mouse buttons simultaneously and drag a rectangle; double-click the graph to return to full size.

Plotting Multiple Selections

You can plot a multiple selection by using standard **Ctrl+click** and **Shift+click** operations.

A legend at the top of the screen displays the key to the graph.

Multiple Plot Outputs




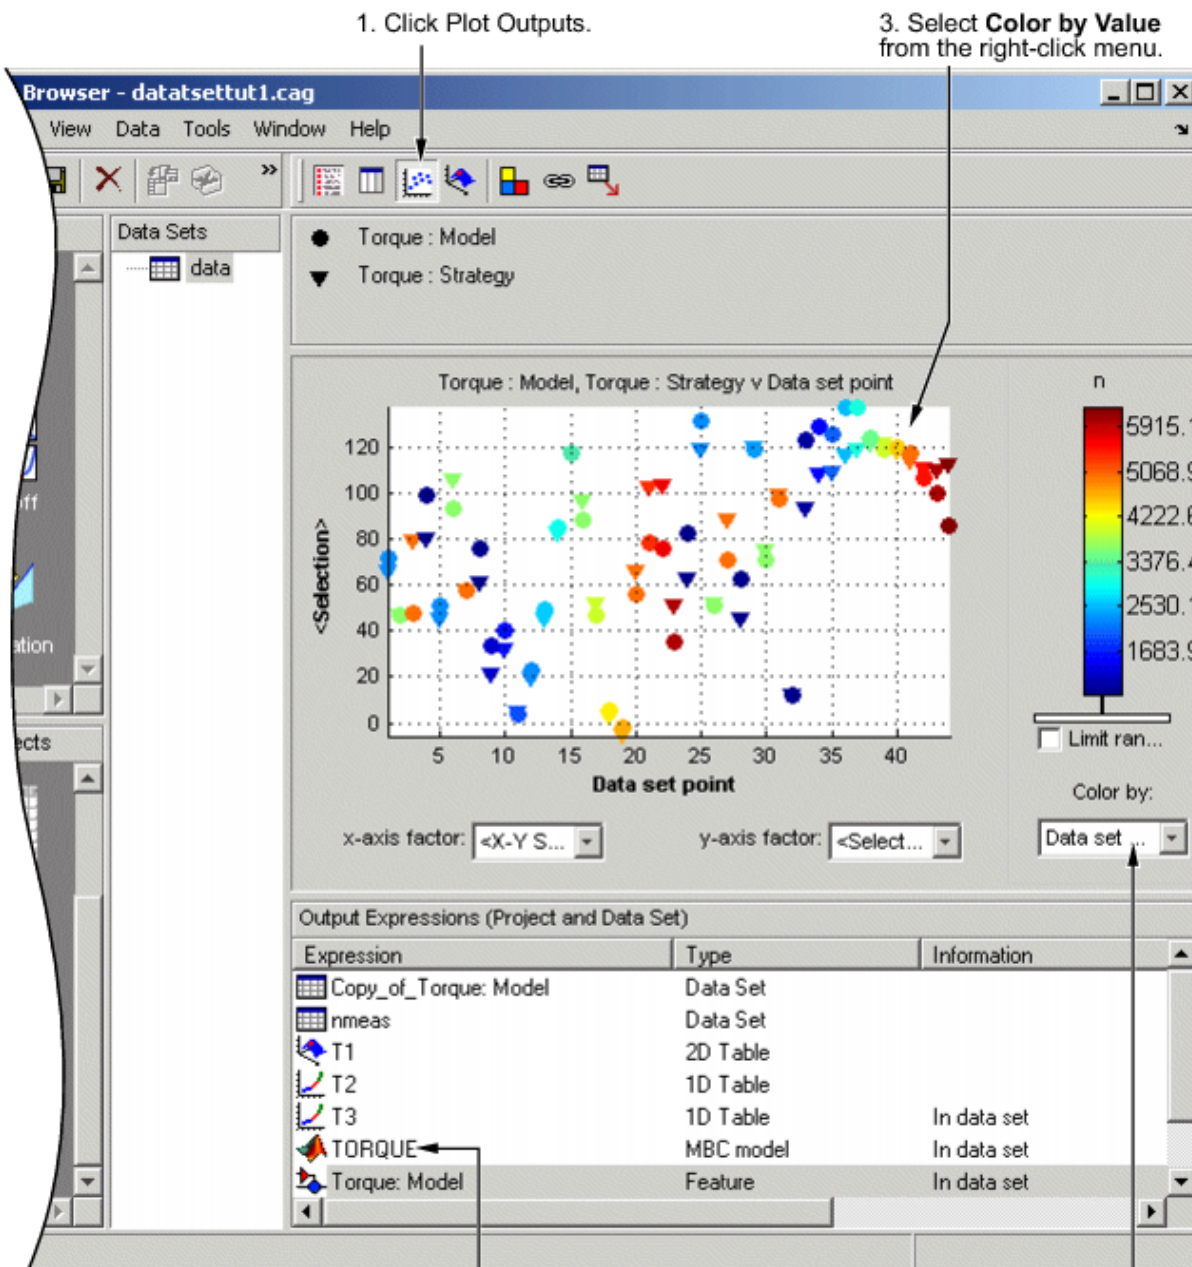
When exactly two items are displayed, further plot options are available:

- Plot the first item against the second item (**X-Y Selection**).
- Display the error using one of the following options:
 - Error
 - Absolute error
 - Relative error (%)
 - Absolute relative error (%)

Use Color to Display Information

You can use the plot view to display more information by coloring the plots.

- 1 Select **View > Plot** or click .
- 2 Highlight the correct expression in the **Output Expressions (Project and Data Set)** pane.
- 3 Select **Color by Value** from the right-click menu of the plot.
- 4 Select from the pop-up menu the variable you want to use to color the plot.



1. Click Plot Outputs.

3. Select **Color by Value** from the right-click menu.

2. Select the expression.

4. Select the correct variable.

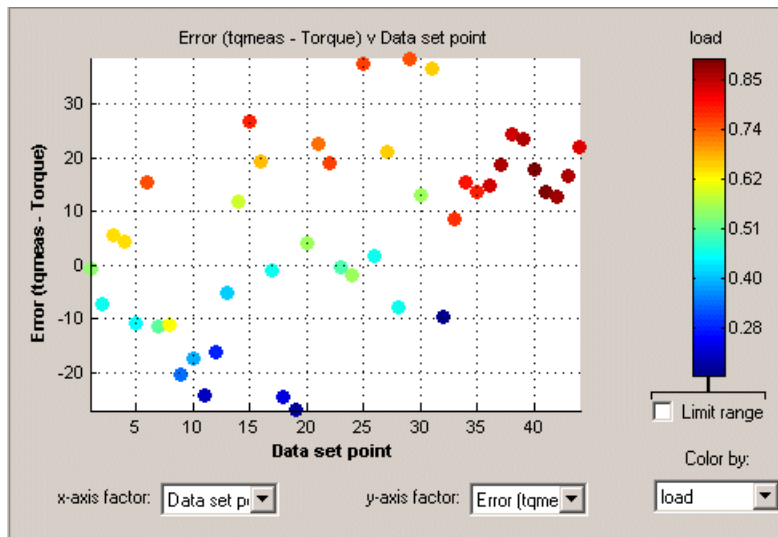
The screenshot shows a software window titled "Browser - datatsettut1.cag" with a menu bar (View, Data, Tools, Window, Help) and a toolbar. The main area displays a scatter plot titled "Torque : Model, Torque : Strategy v Data set point". The plot has a y-axis labeled "<Selection>" ranging from 0 to 120 and an x-axis labeled "Data set point" ranging from 5 to 40. A color scale legend on the right indicates values from 1683.9 (blue) to 5915.1 (red). Below the plot, there are dropdown menus for "x-axis factor" (set to "<X-Y S..."), "y-axis factor" (set to "<Select..."), and "Color by:" (set to "Data set ...").

The "Output Expressions (Project and Data Set)" pane at the bottom contains the following table:

Expression	Type	Information
Copy_of_Torque: Model	Data Set	
nmeas	Data Set	
T1	2D Table	
T2	1D Table	
T3	1D Table	In data set
TORQUE	MBC model	In data set
Torque: Model	Feature	In data set

In the following figure, you can see

- A plot of the Sum vs Data Set Point (this is the strategy from a torque feature calibration).
- The points are colored by load.
- For this example it can be seen that, in general, the higher the load, the higher the value of torque.



Restricting the Color

You might be interested in only part of the display; for example, you might only be interested in points with a low engine speed. The various display options enable you to color only the points that you are interested in.

To restrict the color,


- 1 Select the **Limit range** box, or right-click the plot and select **Limit Color Range**.

Three limit markers appear in the color bar. The colors in the color bar are compressed within the limit markers. This increases the range of colors over the range you are interested in (between the limits), making it easier to see the distribution of points.

- 2 Adjust the maximum, midpoint, and minimum of the range by dragging the limit markers on the color bar.
- 3 Examine the data points and those that are outside the range.

Use the right-click menu to alter the view of the points outside the range:

- Select **Exclude** to remove all points outside the limits from the display.
- Select **Color Outside Limits** to display all points in color, including those outside the limits. Points outside the limits are still colored, but only dark red or dark blue, depending on which end of the range they are.
- Select **No Color Outside Limits** to display the points as in the example shown. Points outside the limits are plotted as empty circles.



Link Factors in a Data Set


A factor can be linked to another. The factor then takes on the values of that other factor, overwriting the original values.


For example, you might want to link a variable spark with a model for maximum brake torque (MBT) to evaluate a torque model.

To link two factors,

- 1 Select **Data -> Links**. This opens a dialog box.
- 2 Select the data set factor that you want to overwrite.

CAGE generates a list of factors that you could possibly link to the selected factor. (For example, you cannot link to a factor that depends on the selected factor.)

- 3 Select the factor that you want to link the selected factor with.
- 4 Click  to link the two factors.




2. Select the factor that you want to overwrite.

4. Click here to link the factors.

2. Select the factor that you want to link it with.

CAGE then overwrites the data set factor with the link.

To break a link and return to normal evaluation, click .

Once all the links have been created or broken as you want, click **OK** to exit the dialog.

See also: "Set Up Data Sets" on page 10-4

Manipulate Models in Data Set View

A model in a data set can be treated as either an input or an output. This is particularly useful when a model is used as an input to another model and you want to view specific values of the input model. For example, linking a model of MBT Spark to a Spark model allows the evaluation of a TQ model at MBT.

To change a model to an input,

- 1 Highlight the desired model in either the factor view or the table view.
- 2 Select **Treat as Input** from the right-click menu.

To revert a model to an output,

- 1 Highlight the desired model in either the factor view or the table view.
- 2 Select **Treat as Output** from the right-click menu.

Export Data Sets

In this section...
“Exporting Data to the Model Browser” on page 10-20
“Exporting Data to File” on page 10-20

Exporting Data to the Model Browser

When viewing a data set, you can export the data to the Model Browser. The Model Browser must be open.

Select **File > Export > Data > Model Browser**.

The Model Browser displays the data set in the Data Sets list at the Project node.


Exporting Data to File

When viewing a data set, you can export the data to a comma-separated value file.

- Select **File > Export > Data > File**.
- In the file browser, specify the file name and location and click **Save**.

Surface Viewer

View a Model or Strategy

To access the Surface Viewer, select **Tools > Surface Viewer** or click  on the toolbar.

These are the main steps to view the model or feature using the **Surface Viewer** dialog box:


- 1 The model or feature selected when you open the **Surface Viewer** is displayed in the plot. If you have more than one model or feature, select what to display from the top **Items** list.

You can multiselect up to four items at once using **Ctrl+click** (the plot view on the right divides into a maximum of four plots). All the settings below the **Items** list apply to all plots. If one of the features selected in the **Items** list does not contain the appropriate input variables you select to plot, there is no plot for that item.

- 2 Select the ranges for the variables.
- 3 Choose the plot type to display. You can view surfaces, contour plots, single and multilines, movies, tables, and single values.

For example, as you view a feature, you can view either the strategy, the model associated with that feature, the error between the model and the strategy, or the prediction error if the model was imported from the Model Browser. You can also use one of these factors to shade the surface formed by one of the other factors, and you can select any two factors to display simultaneously as two surfaces.

- You can make a movie. This enables you to view the model or feature as it steps through several values of a variable. For example, if you want to view a feature calibrated for maximum brake torque (MBT) as it varies over exhaust gas recycling (EGR), you can make a movie of the feature.
- You can also print or export the display.



See Also


More About

- “Display Errors” on page 11-12
- “Display the Model or Feature” on page 11-6
- “Print and Export the Display” on page 11-13
- “Set Variable Ranges” on page 11-4
- “View a Model or Strategy” on page 11-2

Set Variable Ranges

The **Surface Viewer** does not work over continuous ranges, only at discrete points. Specify, for the model or feature, the discrete points you want to include in the display. You can display models or features over a range of points. To edit the displayed values of a variable, double-click in the value box for the appropriate variable.

- Variables not being used for the axes plotted have a single value for that plot; to edit the displayed value for these variables you can type directly into the edit box after double-clicking.
- For variables specified by the axes drop-down menus, the value box displays the range over which that variable is plotted and the number of points plotted across that range. To edit both the range and the number of points, double-click the value box. The **Value Editor** opens.



Here you can indicate the points to include in the display. You can specify

- The minimum and maximum values and the number of points across that range by choosing **Uniform Vector** and typing in the edit boxes **Min**, **Max**, and **Number of points**.
- Each discrete point at which you want to evaluate the model (or feature), by choosing **Freeform vector**, and then typing the required values.

For example, if you want to display the variable x at 0, 1, 7, 30, and 50, enter the following in the **Freeform vector** edit box, separated by tabs or spaces:

```
0 1 7 30 50
```

Click **OK** to apply your changes to the plot.

When you alter the variables, you can select whether you want the display to update automatically or not. You can toggle the automatic update on and off by selecting **Tools > Auto-Evaluate**. When you want to update the display, select **Tools > Evaluate Now**. Both of these options have equivalent toolbar buttons:




Display Point-by-Point Models in the Surface Viewer

When you are displaying a point-by-point model, you can select the operating point to display. When you are using point-by-point models, these are the points of interest you want to display.

To select the operating point to display in the Surface Viewer,

- 1 Select **Tools > Select Operating Point** (or the equivalent toolbar button). The Model Operating Points dialog box opens.



- 2 Select the operating point you want to display and click **OK**.

Surface Viewer snaps the display automatically to the selected point-by-point model operating point. When you select an operating point, Surface Viewer uses the model ranges for that operating point to set the local inputs (ranges and midpoints as applicable).

See Also


More About

- “Display Errors” on page 11-12
- “Print and Export the Display” on page 11-13
- “View a Model or Strategy” on page 11-2

Display the Model or Feature

Use Display Options

The **Plot Type** drop-down menu gives the options on how to display the model or feature.



Use the options in this menu to display the model or feature as described in the following sections.


When plotting multiple models or features, it is useful to link axes rotation or use common Y- or Z-ranges. Use the display options (toolbar button or **View** menu).

In any of these views you can select **View > Statistics**, or click the equivalent toolbar button. This opens a dialog box with a list of the summary statistics (mean, standard deviation, maximum, or minimum) of your currently selected model, strategy, or error for the current display.

For the plots (not movie, single value or tables) you can use the **File** menu or toolbar to print, copy to clipboard or print to figure. You can also export plot values to CSV file.

You can alter display options for all plots (not value or tables) with the **View** menu or toolbar button.

Surface



You can rotate the surface plots by left-clicking and dragging.

If you are using the surface viewer to view a feature, you can choose the following options to display:

- Model
- Strategy
- Prediction Error
- Error (between the model and the strategy)

When viewing models, there are no strategy options. You can choose these options from the drop-down menus for **Surface 1 Height**, **Surface 1 Shading**, and **Surface 2 Height**.

The image shows a dialog box titled "Data to Plot". It contains the following controls:

- Plot type:** A drop-down menu with "Surface" selected.
- Surface 1 Height :** A drop-down menu with "Model" selected.
- Surface 1 Shading :** A drop-down menu with "None" selected.
- Surface 2 Height :** A drop-down menu with "None" selected.
- Mark points outside boundary**

You can view any of these options alone as a primary surface (by leaving the last two options set to **None**). You can add a second option to shade the primary surface, for example to color your model surface with the error between the model and the strategy, to highlight problem areas.


When you choose to shade a primary surface, a color bar appears to the right of the plot to show you the scale. You can change the maximum and minimum values of the shading factor by typing in the edit boxes above and below the color bar.

You can add a second surface to display any two of the options simultaneously, for example, your model and your strategy.

If you have a boundary model, you can display the boundary by selecting the check box.

Select the **Inputs** to plot from the **X-axis** and **Y-axis** drop-down lists, and specify the ranges of inputs in **Value** controls.


Contour



You can specify where you want contours by clicking **Set Contour Values**. Use the check box to return to automatic contour value selection. You can also control number of contours, filling, and labels in the display options (toolbar or **View** menu).

You can enable **Cursor Mode** (use the **View** menu or toolbar button) and then click the plot lines to display the values at a point (plotted with an X). The values are shown in the status bar.

Line




A line plot - you can display up to three different lines (strategy, model, prediction error, and error between the model and strategy). Use the **Line** drop-down lists to select what to plot. You can select the check box to clip to a boundary if available.

You can enable **Cursor Mode** (use the **View** menu or toolbar button) and then click the plot lines to display the values at a point (plotted with an X). The values are shown in the status bar.

Single Value

This displays the value of the model, strategy, prediction error, or error at the point you have specified in the variable value boxes.

Multiline



Select the variables to plot from the **X-axis** and **Line colors** drop-down menus. Control the number of lines by altering the **Values**. You can use the check box to clip to a boundary if available.

You can enable **Cursor Mode** (use the **View** menu or toolbar button) and then click the plot lines to display the values at a point (plotted with an X). The values are shown in the status bar.

Table

Project/Branch 1/Fn_Feature			
x \ y	0.000	0.500	1.000
-5.000	35.000	33.776	30.403
-4.500	30.250	29.026	25.653
-4.000	26.000	24.776	21.403
-3.500	22.250	21.026	17.653
-3.000	19.000	17.776	14.403
-2.500	16.250	15.026	11.653
-2.000	14.000	12.776	9.403
-1.500	12.250	11.026	7.653
-1.000	11.000	9.776	6.403
-0.500	10.250	9.026	5.653
0.000	10.000	8.776	5.403
0.500	10.250	9.026	5.653
1.000	11.000	9.776	6.403
1.500	12.250	11.026	7.653

You can select a 2-D or 1-D table to display. Select the check box to mark cells outside the boundary.

Choose variables for the axes of your table and set the range and number of points. Set single values for any other variables.

For 2-D tables, you can use the **Cell values** drop-down menu to select whether to display the model output or the prediction error.

For 1-D tables you can select what to display in columns 1, 2 and 3: Model, Prediction error, Strategy or Error (strategy – model) (for features), or choose None for 2 and 3 to display only a single column. When viewing models, there are no strategy options.


See Also

More About

- “Display Errors” on page 11-12
- “Print and Export the Display” on page 11-13
- “Set Variable Ranges” on page 11-4
- “View a Model or Strategy” on page 11-2

Make Movies

Choose **Movie** from the **Plot Type** drop-down menu in the **Data to Plot** pane.



The movie option allows you to see an evaluation over two variables at successive values of a third variable. For example, a model of torque might have speed (N), load (L), and air/fuel ratio (A) as inputs.

The movie option allows you to view how the torque model behaves over the ranges of speed and load for successive values of air/fuel ratio.

- 1 Select three variables from the **X-axis**, **Y-axis**, and **Time** drop-down menus, to indicate which variable you want to display. You can view the model surface plotted across the range of two variables, and define the third variable as "time" to see the model surface change across the third variable range.
- 2 Define the variable ranges using the **Value** boxes for the inputs.
- 3 Select the check box to mark boundaries if available.
- 4 Click **Play**.
- 5 You can click the buttons at each end of the progress bar under the plot to step through the movie, or click anywhere along the bar (or click and drag the blue pointer) to display a particular point in the movie. You can rotate the plot (including during play).

See Also

More About

- "Set Variable Ranges" on page 11-4

Display Errors

Introducing Error Displays

There are two different error displays available in the surface display options for primary and secondary surfaces and surface shading:

- Error between the model and the strategy
- Prediction error of the model

Feature Error Data

When you are viewing a feature, this displays the error between the strategy and the model.

To display the error, select **Error (strategy-model)** from the drop-down menu for primary or secondary surface. You can also choose to shade your primary surface with the error by using the **Surface 1 Shading** menu.

To view the error statistics, select **View > Statistics**. This opens a dialog box with a list of the summary statistics for the error between model or feature.

Prediction Error Data

If the model is imported from the Model Browser, it is possible to display the *prediction error* (PE) data.

Prediction Error Variance (PEV) is a useful way to investigate the predictive capability of your model. It gives a measure of the precision of a model prediction. You can examine the PEV in the Model Browser, both in the **Prediction Error Variance Viewer** and to shade surfaces in the **Model Selection** and **Model Evaluation** views. Here you can examine the PEV of designs and models. When you export the model to CAGE, you can see this data in the **Surface Viewer** in the Prediction Error option. See the Model Browser GUI Reference and Technical Documents for details about the calculation of Prediction Error.

Viewing the Prediction Error

Select Prediction Error from the drop-down display menus for primary or secondary surfaces. You can also choose Prediction Error to shade your primary surface. As with all other plots, you can view the statistics for the Prediction Error displayed by selecting **View > Statistics**. The mean, standard deviation, and so on, are calculated over the range specified in the variable value boxes.

See Also

More About

- “Display the Model or Feature” on page 11-6
- “Print and Export the Display” on page 11-13
- “View a Model or Strategy” on page 11-2

Print and Export the Display

To print the display, select **File -> Print**, or you can select Print to Figure. Selecting **File > Copy to Clipboard** copies the plot image to the clipboard. This is useful if you want to place plot images into other applications. These print options also have equivalent toolbar buttons.

You can also export the display data to a comma-separated variable file.

To export the display, select **File > Export to CSV**. The currently selected option is exported. The primary input to the first plot is exported (this is the top left if you have multiple plots). The output is the values at the grid of points specified by the current ranges and input values. The inputs for shading and secondary surfaces are not exported.

You cannot print table plots, but you can click and drag to select cells and press **Ctrl-C** to copy the values to the clipboard, or you can export them to CSV files and then load them into Excel.

See Also

More About

- “Display Errors” on page 11-12
- “Display the Model or Feature” on page 11-6
- “Set Variable Ranges” on page 11-4
- “View a Model or Strategy” on page 11-2

